



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

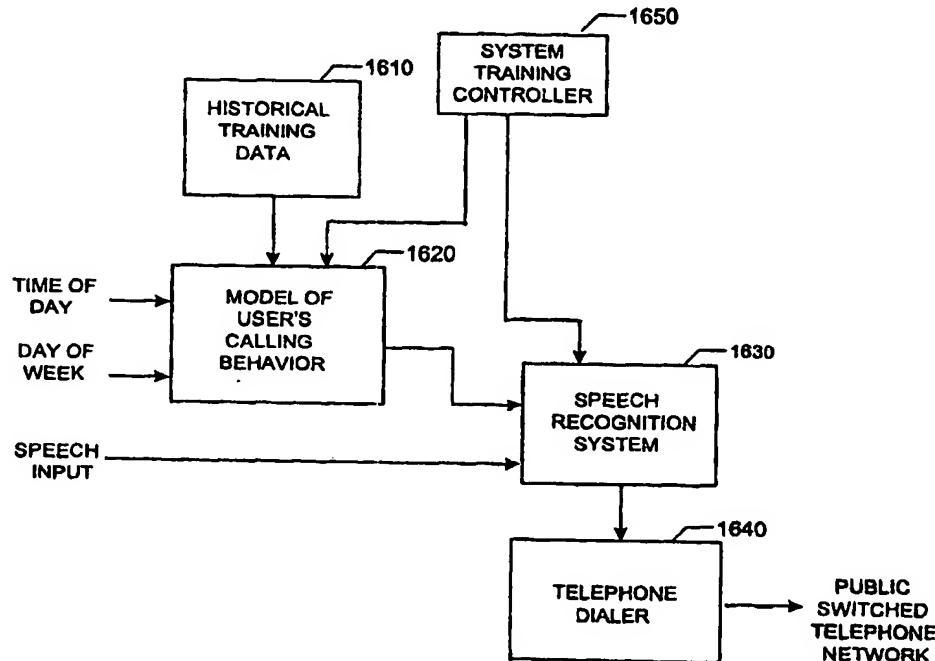
(51) International Patent Classification ⁶ : H04M 1/27	A1	(11) International Publication Number: WO 98/16048
		(43) International Publication Date: 16 April 1998 (16.04.98)

(21) International Application Number: PCT/US97/17623	(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).
(22) International Filing Date: 6 October 1997 (06.10.97)	
(30) Priority Data: 08/726,604 7 October 1996 (07.10.96) US 08/942,201 1 October 1997 (01.10.97) US	
(71) Applicant: NORTHERN TELECOM LIMITED [CA/CA]; Station A, P.O. Box 6123, Montreal, Quebec H3C 3J5 (CA).	
(71)(72) Applicant and Inventor: WILL, Craig, Alexander [US/US]; 37675 Fremont Boulevard No. 23, Fremont, CA 94536 (US).	Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>
(74) Agents: GARRETT, Arthur, S. et al.; Finnegan, Henderson, Farabow, Garrett & Dunner, L.L.P., 1300 I Street, N.W., Washington, DC 20005-3315 (US).	

(54) Title: VOICE-DIALING SYSTEM USING MODEL OF CALLING BEHAVIOR

(57) Abstract

A method and apparatus for assisting voice-dialing using a model of an individual's calling behavior to improve recognition of an input name corresponding to a desired telephone number. When the individual picks up a telephone, activity is initiated in a neural network model of the individual's calling behavior that predicts the likelihood that different numbers will be called, given such predictors as the day of the week and the time of day. The model is constructed by training the neural network with data from the user's history of making and receiving telephone calls. The auditory output from an automatic speech recognition system and the output from the user model



are integrated together so as to select the number that is most likely to be the number desired by the speaker. The system can also provide automatic directory assistance, by speaking the number aloud rather than dialing it. In one version, the system is a personal directory for an individual maintained on that individual's personal computer. In another version, the system serves as a directory for a given physical or virtual site, with information about the institutional organization at the site in addition to individual calling histories used to track calling patterns and make predictions about the likelihood of calls within the site.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Republic of Korea	RO	Romania		
CU	Cuba	LC	Kazakhstan	RU	Russian Federation		
CZ	Czech Republic	LI	Saint Lucia	SD	Sudan		
DE	Germany	LK	Liechtenstein	SE	Sweden		
DK	Denmark	LR	Sri Lanka	SG	Singapore		
EE	Estonia		Liberia				

VOICE-DIALING SYSTEM USING MODEL OF CALLING BEHAVIOR

BACKGROUND OF THE INVENTIONA. Field of the Invention

This invention relates generally to systems for telephonic communications with audio message storage and retrieval and, more particularly, to telephonic communications involving repertory or abbreviated call signal generation and abbreviated dialing. The invention further relates to systems based on artificial intelligence techniques, particularly those using knowledge processing, and especially to adaptive or trainable systems that create sets of rules and use parallel distributed processing components.

10 B. Description of the Related Art

Both rotary and touch-tone dialing rely on telephone numbers to initiate desired telephone connections. Telephone companies use the numbers to route calls, but people now depend on the numbers for all telephone communications. This is somewhat unnatural because people generally select those with whom they would like 15 to talk by name or other convention. Indeed, telephone directories are arranged by name, not number.

Some companies started to develop voice-activated dialing systems to replace touch-tone dialing. In such systems, telephone users speak the name of an individual or destination into the microphone of a telephone handset to initiate a telephone call.

20 Voice-dialing thus allows connection to be made directly, avoiding the step of looking up names to locate corresponding telephone numbers.

Examples of experimental voice-dialing systems appear in L. R. Rabiner, J. G.

Wilpon, and A. E. Rosenberg, "A voice-controlled, repertory-dialer system," Bell System Technical Journal, Vol. 59, No. 7 (September, 1980), and U.S. Patent No. 4,348,550 to Pirz et al. These systems have limited accuracy and speed and cost a great deal of money.

5 Recent advances in speech recognition have improved performance dramatically, particularly for systems that are not trained to a particular speaker that have, until recently, performed worse than systems trained to particular speakers. In addition, the increasing computational and memory capacity and decreasing cost of computing hardware improve the commercial viability for simpler applications of
10 speech recognition such as voice-dialing.

Limitations on the performance of voice-dialing systems, however, still significantly reduce their commercial applicability. Such systems frequently make errors, with the rate of errors increasing with increased vocabulary size and factors such as environmental noise, unusual accents, and the use of foreign or unusual
15 names that are difficult to pronounce consistently. The limited accuracy of recognition performance resulting from these factors restricts the possible range of applications for conventional voice-dialing systems by limiting the vocabulary, environment, user population, and hardware platforms on which the systems can run.

It is therefore desirable to seek techniques that will improve the accuracy and
20 speed of speech recognition performance in voice-dialing systems. A number of alternative techniques have been used in the past. One approach is to ask the user for verification before dialing ("Did you say Anatoly Korsakov?"), and presenting a

different name if the user says "No." See, for example, U.S. Patent No. 5,222,121 to Shimada. Another approach, disclosed by Iida et al. (U.S. Patent No. 4,864,622), modifies or replaces the voice template used for speech recognition when the template is not performing adequately.

5 None of these approaches, however, really improves speech recognition performance for voice-dialing systems. They merely require additional user interaction to assist in the voice-dialing process.

SUMMARY OF THE INVENTION

There is, therefore, a need to improve the speed and accuracy of voice-dialing
10 systems. There is also a related need to allow such systems to adapt and learn.

The present invention meets these needs using a neural network that creates a model of the telephone calling behavior of an individual and uses this model to increase the performance of the automatic speech recognition system that matches incoming spoken names with names stored in a directory.

15 To achieve the objects and in accordance with the purpose of the invention, as embodied and broadly described herein, provides a method and apparatus for assisting voice-dialing by receiving voice input from a user representing a name corresponding to a desired telephone number, selecting stored names that most closely match the voice input, predicting a likelihood of the user calling telephone numbers based on a
20 model of the user's calling behavior, and determining the desired telephone number according to the predicted likelihood of the user calling the telephone number corresponding to each selected name.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate preferred embodiments of the invention and, together with the description, explain the goals, advantages and principles of the invention. In 5 the drawings,

FIG. 1 is a block diagram of hardware architecture according to a preferred embodiment of the voice-dialing system of the present invention;

FIG. 2 is a functional flowchart of the process steps used to initiate telephone calls according to the preferred embodiment of the voice-dialing system of the 10 present invention;

FIG. 3 is a block diagram of the software components according to the preferred embodiment of the voice-dialing system of the present invention;

FIG. 4 is a diagram used to explain the architecture of a neural network that models the user's calling behavior for the preferred embodiment of the voice-dialing 15 system of the present invention;

FIG. 5 is a diagram used to explain the architecture of an integrator neural network for the preferred embodiment of the voice-dialing system of the present invention;

FIGs. 6a and 6b show a functional flowchart of steps used by the voice-dialing 20 system in FIG. 3, during incoming and outgoing telephone calls, to record information for training the neural networks shown in FIGs. 4 and 5;

FIG. 7 shows the data structure of historical call information used for training

the calling behavior neural network according to the preferred embodiment of the voice-dialing system of the present invention;

FIG. 8 shows the data structure of historical call information used for training the integrator neural network according to the preferred embodiment of the voice-dialing system of the present invention;

5

FIG. 9 is a flowchart of events that occur when the preferred embodiment of the voice-dialing system of the present invention trains the neural networks shown in FIGs. 4 and 5;

10 FIG. 10 is a flowchart of the steps used when the preferred embodiment of the voice-dialing system of the present invention trains the calling behavior neural network;

FIG. 11 is a flowchart of the steps used when the preferred embodiment of the voice-dialing system of the present invention trains the integrator neural network;

15 FIG. 12 is a flowchart showing the procedure followed by the preferred embodiment of the voice-dialing system of the present invention when the user modifies the directory of names and associated telephone numbers;

FIG. 13 is a block diagram of another embodiment of the voice-dialing system according to the present invention;

20 FIG. 14 is a block diagram of the software components for the system in FIG. 13;

FIG. 15 is a diagram used to explain the architecture of a category-based calling behavior neural network for the system in FIG. 13; and

FIG. 16 shows a block diagram of an alternative architecture for the voice-dialing system that was previously shown in FIG. 3.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Reference will now be made in detail to the preferred implementation of the present invention as illustrated in the accompanying drawings. Wherever possible, 5 the same reference numbers will be used throughout the drawings and the following description to refer to the same or like parts.

A voice-activated dialing system according to the present invention is built around a personal directory stored in the memory of a personal computer that holds 10 names and associated telephone numbers. The system can be used either locally, by picking up a telephone and speaking the name associated with the desired number, or by connecting from a remote location and speaking the name. It may be implemented in a personal computer that is provided with a telephone interface card, as well as 15 software to perform speech recognition and speech synthesis, to implement a neural network and dial a telephone number, and to control the voice-dialing system. It may also be used provide automatic directory assistance by speaking the number aloud rather than dialing it.

The architecture of the system consists of three components: a component that processes incoming speech and matches it against representations of the names in the 20 personal directory, a component that models the user's calling behavior, and a component that integrates the outputs of the first two components to produce the name that the user most likely desires to call.

The user calling behavior model component consists of a multilayer feedforward neural network that uses the backward propagation learning algorithm.

The inputs to the neural network accept the current date and time, while the output of the network provides a signal for each telephone number in the directory. The speech 5 recognition component of the system processes an auditory input and a stored list of names in either a textual or auditory representation to provide a list of those names that best match the auditory signal and a measure of the quality of each match.

The component of the system that integrates the outputs of the first two components also consists of a multilayer feedforward neural network using backward 10 propagation. The inputs to this neural network include one input for each telephone number in the directory from the output of the calling behavior model network, and one input for each telephone number from the output of the speech recognizer.

According to another aspect of the present invention, a voice-activated dialing system consists of a microprocessor-based server for a PBX system that implements a 15 voice-dialing directory for a given physical or virtual site. The voice-dialing system makes use of three neural networks for a given individual, including the user calling behavior model and the integrator neural network. The third neural network is common to all individuals at the site, and implements a predictive model of calling between individuals at the site. This neural network is a multilayer feedforward 20 neural network that uses the backward propagation learning algorithm. Every telephone number at the site corresponds to a category, with the category assignment made according to the structure of the organization at the site. The common network

contains an input unit for each category and an output unit for each category.

A. Personal Directory System

1. Hardware Architecture

FIG. 1 shows the hardware architecture for a preferred embodiment of the

5 voice-dialing system according to the present invention implemented as a personal directory system for an individual. Personal directory system 100 includes a workstation 110, which includes hardware for a standard personal computer (for example, an IBM compatible personal computer), together with some additions related to telephony, and an ordinary telephone 120 (for example, a touch-tone telephone). Alternatively, telephone 120 may be connected to workstation 110 when 10 workstation 110 includes required voice input and output devices (not shown) that perform functions comparable to telephone 120.

Workstation 110 consists of microprocessor 140, random access memory (RAM) 150, hard disk 160, floppy disk and drive 170, video display 180, keyboard 15 190, and mouse 195. These may be standard off-the-shelf hardware. For example, microprocessor 140 may be a Pentium® processor manufactured by Intel Corp., USA, and video display 180 may be a NEC MultiSync 3V monitor manufactured by NEC Corp., Japan.

System 100 also includes telephone port 130 connecting the workstation 110 20 to a public switched telephone network. Alternatively, workstation 110 may be connected directly to a PBX via a digital connection for both voice and control signaling.

Telephone port 130 includes a switch, controlled by microprocessor 140 and also by DTMF tone receivers in the telephone port 130, to connect telephone 120 to the public switched telephone network or to microprocessor 140. Microprocessor 140 can also be connected directly to the public switched telephone network to allow 5 dialing a number for an individual user. Telephone port 130 for use in a home or small office environment includes analog to digital and digital to analog converters and mechanisms to receive and transmit DTMF codes either via specialized hardware or with software.

All processing for the voice-dialing system may be done with microprocessor 10 140 as FIG. 1 shows. Workstation 110 may also include one or more specialized digital signal processing chips as coprocessors for linear predictive coding for speech recognition, format synthesis for speech synthesis, or processing and learning for a neural network.

2. Operation Overview

15 FIG. 2 shows a flowchart of a voice-dialing procedure 200. The steps of procedure 200 are implemented in software and use personal directory system 100 to control voice-dialing. The software, which may be stored in RAM 150, is executed by microprocessor 140.

20 The flow chart assumes that a user has previously created a database of names and associated telephone numbers. The database may be stored on hard disk 160. One such conventional software package that may be used to create such a database is Microsoft Schedule+®, manufactured by Microsoft Corporation.

The software for the voice-dialing procedure 200 runs as a background process on workstation 110, and microprocessor 140 periodically tests whether the handset of the telephone instrument is off-hook (step 205). When it becomes off-hook, microprocessor 140 activates processing on a previously constructed neural network related to the user's calling behavior. The calling behavior neural network 5 predicts the likelihood that a user will make a call to each number in the database, given the history of calling behavior encoded in the model and the current time of day and day of the week the new call is being made (step 210).

When the calling behavior neural network is activated, microprocessor 140 10 may also play an auditory signal to the user. Such a signal is not necessary for all implementations, however, to indicate that the system is ready for use.

The user then speaks the name associated with a desired number and microprocessor 140 tests whether the user has spoken (step 215). If so, microprocessor 140 processes the speech to extract the appropriate features and 15 matches the results against the names in the database to find the best matches (step 220). If the spoken input does not match any name above a certain minimum threshold of similarity, test recognition fails (step 225), and microprocessor 140 awaits further spoken input (step 215).

If the spoken input matches at least one name, microprocessor 140 combines a 20 similarity measure from the speech recognizer for each match with the resulting likelihood for the corresponding number from the calling behavior model to determine the name and number the user most likely intends to call (step 230).

Microprocessor 140 then plays a name in auditory form to the user through the handset (step 235), and the user signals his or her agreement by responding verbally with either "Yes" or "No" (step 240). If the user responds with "Yes" (step 240), microprocessor 140 retrieves the number and dials it (step 245). Microprocessor 140 5 also saves the transaction in a training database located on hard disk 160 (step 250).

If the user responds with "No" (step 240), microprocessor 140 determines the next best overall matching name and number (step 255), and plays it to the user (step 235). If there is no adequate next best name (step 255), a recorded voice asks the user "Who do you want to call?" (step 260) and control passes to allow the user to try 10 again to speak the desired name (step 215). The test of adequacy can be based on either a fixed number, e.g., 3, of names provided to the user, a minimum threshold for quality of the match, or a combination of both. This process continues until a user verifies a name or hangs up.

After dialing a number, microprocessor 140 periodically tests to see if the user 15 has hung up the handset (step 265). If so, microprocessor 140 monitors the handset for an off-hook condition to initiate another call (step 205).

A "No" response to a name spoken to the user for verification (steps 235, 240) can also cause microprocessor 140 to save a record in the training database for either 20 the integrator neural network, the call behavior model neural network, or both.

3. Software Components

FIG. 3 shows a block diagram of the software system 300 executed by microprocessor 140. The software system 300 may be stored on hard disk 160.

System 300 consists of three primary components: a model of the user's calling behavior 320, a speech recognition system 330 that performs automatic speech recognition, and an integrator 350 that integrates the outputs of the first two components to produce the best estimate of the name the user desires to call. System 5 300 also includes a telephone dialer 360 that looks up the actual telephone number in a table and dials it.

System training controller 370 trains the calling behavior model 320 and integrator 350, using historical training data 310 and 340, respectively. System training controller 370 is described in detail below with reference to FIGs. 9-11.

10 Both calling behavior model 320 and integrator component 350 preferably include a neural network. These neural networks use historical training data 310 and 340, respectively, that are maintained to continue training the neural networks when appropriate.

15 The use of separate neural networks for modeling calling behavior and integration reduces the complexity of the voice-dialing system and allows separate training of each network.

When a user picks up the handset of telephone 120 or dials in to the workstation 110 from a remote telephone and identifies himself or herself, microprocessor 140 reads in the weights of the user's calling behavior model 320 20 from hard disk 160 and determines the current time and day of the week. When the user speaks the name of the person to be called, speech recognition system 330 processes the input speech data and attempts to match it against the set of stored

representations, typically sequences of phonemes, that represent each name in the database. An example of a speech recognition system with the desired capabilities include the "Model asr1500/M" speech engine from Lernout & Hauspie Speech Products N.V., Ieper, Belgium. These systems run on a personal computer with a 5 Pentium® microprocessor in close to real time without needing an additional coprocessor.

Speech recognition system 330 produces sets of floating point numbers, each representing the extent to which there is a match between the speech input and the stored representation for the name associated with each telephone number. In 10 practice, commercially available speech recognition engines typically produce an output consisting of a list of the "N best" matches to names in the database for which the match was above a given threshold value, with a quality measure for each. The quality measure for all other items in memory can be regarded as 0.

Integrator 350 receives the output data from both the user's calling behavior 15 model 320 and speech recognition system and produces an output consisting of the best telephone numbers by applying the inputs to integrator 350's own neural network. This number may be dialed immediately, or a protocol followed that asks the user to verify the number as correct before dialing it (see FIG. 2). The call 20 attempt itself is recorded in a historical training database and stored on hard disk 160 so that it can be used as historical training data to train the user's calling behavior model 320.

Integrator 350 can be implemented by conventional techniques. One such

approach is simply to determine weightings that indicate the relative contribution of the calling behavior model 320 and the output of speech recognition system 330 to making the best prediction of the number the user intended to call. The output from the user's calling behavior model 320 and speech recognition system 330 for each candidate number is multiplied by the given weight and then summed, and the number with the highest numerical score then selected. This embodiment is simpler and reduces the computational requirements of the system in FIG. 3. The weighting would be arbitrarily fixed, however, and would not be adjusted specifically for each name and number and would thus be less accurate.

10 a. User's Calling Behavior Model Neural Network

FIG. 4 shows the architecture of a neural network 400 that models the user's calling behavior for the voice-dialing system 300. Network 400 is shown as a three-layer feedforward neural network, and consists of an input layer 410, a hidden layer 420, and an output layer 430. Such a network architecture is described in detail in the paper by D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning internal representations by error propagation," Parallel Distributed Processing: Explorations in the Microstructure of Cognition, J. E. McClelland, D. E. Rumelhart, and the PDP Research Group, Editors, Cambridge, MA: MIT Press, 1986, Vol. 1, pp. 318-362.

15 Mathematical equations that describe the computation of the activity level of a unit from its inputs and the role of the weights of connections in such computations can be found in the paper by Rumelhart, Hinton, and Williams as well as in textbooks on neural network architectures and applications.

Network 400 is implemented by software and input values are set to 0.0 for false and 1.0 for true. The software to simulate network 400 is implemented in the C++ programming language and developed in the Microsoft Visual C++® programming environment, including Microsoft Developer Studio® and Microsoft

5 Foundation Class®, all running under the Windows 95® or Windows NT® operating systems.

A neural network consists of "units," or artificial neurons, that receive inputs through "connections" from other units that are essentially artificial resistors. Each such connection has a value known as a weight that is analogous to the resistance of a

10 resistor. Each unit sums the input signal values received from its inputs after being weighted by the connection, and then applies a nonlinear mathematical function to determine a value known as the "activity level" for that unit. This activity level is then provided, after processing it through an output function as the output of the unit and then applied, through the resistive connections, to units in the next highest layer.

15 For example, the outputs of layer 410 are inputs to layer 420.

Input layer 410, with its input units, is actually a dummy layer in which the activity level for each input unit is simply set to the analog value provided as input to each unit. Each input unit is connected to the input of every unit in hidden layer 420. The large arrow 425 represents such full connections.

20 There are approximately as many units in the hidden layer 420 as there are telephone numbers in the directory of the user. Units in layer 420 are called "hidden units" because their values are not directly observable, unlike the units of input layer

410 and output layer 430. The output of each unit in hidden layer 420 is connected to the input of every unit in output layer 430.

The output of each output unit is provided to the rest of the system as the output of neural network 400. In a feedforward network, the flow of information in network 400 is in one direction only, from input layer 410 to hidden layer 420 and from hidden layer 420 to output layer 430, as arrows 425 and 435 show.

When information is applied to the input of network 400, it propagates to hidden layer 420 and then to output layer 430. The value of each output unit, for which there is one unit corresponding to each number in the user's telephone directory, represents the likelihood that that number will be the next number called by the user.

10 Input layer 410 consists of two groups of inputs 413 and 416. First group 413 encodes the current day of the week and consists of 7 units, one for each day of the week. Second group 416 encodes the current time of day and consists of 7 units, each 15 indicating a time within one of the following seven categories: midnight-6am, 6-9am, 9-12 am, 12-1, 1-4 pm, 4-6 pm, and 6-12 pm.

Calling behavior component 320 first determines the current day and time by means of the appropriate systems call, such as GetLocalTime, a calling program in C++, and then codes this information by selecting the appropriate inputs.

20 The day of the week and time of day inputs are not the most significant effects on calling behavior network 400 that result in an output. In most cases, network 400's most significant predictive capability comes from biasing toward or against specific

numbers. Biassing results from training network 400 from the historical data that is relatively independent of day and time. Day and time inputs become significant primarily when very strong patterns occur involving these parameters, such as making many calls to a particular number on the same day and time. Whether inputs exist is not even critical to the operation of the network. User model 400 could reliably predict the likelihood of calls to particular numbers based on the historical training data alone without any inputs to the network. This is because network 400 bases its predictions on a user's calling behavior determined by the frequency of incoming and outgoing calls.

10 Alternatively, neural network 400 may include two layers of hidden units. The additional hidden layer requires an additional set of connections and weights. Each of the two layers has approximately the same number of hidden units, which approximates the number of telephone numbers in the user's personal directory. The advantages of the additional layer are to allow the capture of more subtle interactions among specific numbers, times, and days, than is possible with a single hidden layer.

15 The disadvantages includes additional processing capacity and memory required to implement the network, longer training times, and possibly less stable training.

 The two possible methods for training the neural network 400 are complete and incremental. Complete training is preferred, but it takes place only once per day.

20 If network 400 is not immediately updated for calls made or received during each day, there may be a drop off in accuracy due to a risk of possible data loss. To accomodate for this potential data loss, calls received and made on a particular day

may be kept in RAM 150, with the predicted probability for the telephone number of such a call calculated by a simple procedural algorithm. Calculation of the likelihood of a telephone number corresponding to a call in RAM 150 may be done by simply setting the likelihood to 0.9, and ignoring the prediction made by the network 400.

5 Otherwise, the prediction made by the network 400 would be used.

Incremental training is done after each call whenever it appears that the computer is not being heavily used and computational capacity is available, and consists of that additional training necessary to update network 400 to the just completed call or calls.

10 b. Integrator Neural Network

FIG. 5 illustrates the architecture of integrator 350's neural network 500. The network 500 consists of a multilayer feedforward network with an input layer 510, a hidden layer 520, and an output layer 530. Input layer 510 consists of two groups 515 and 516. First group 515 consists of an input unit for each telephone number defined in the user's directory, with the input connected to the corresponding output for that telephone number from neural network 400. Second group 516 also consists of an input unit for each defined telephone number, with the input connected to the corresponding output for that telephone number from the speech recognition system. The input telephone numbers for which there is an "N best" recognition output for the corresponding name from the speech recognition system have the appropriately scaled (0.0 to 1.0 range) similarity measure fed into the corresponding input units.

15 Telephone numbers for which there is no recognition output for the corresponding

20

name have the corresponding input unit set to 0.0. Network 500 also has hidden layer 520, with the number of units approximating the number of telephone numbers in the directory of the user, and output layer 530, which has one unit for each telephone number in the database.

5 Network 500 can be implemented by software. Once input data is provided to the speech input units, the simulator computes the activity levels for each hidden unit based on all input units, including day and time, and then the activity levels for each output unit. When the information has completed propagating through network 500 in this way, system 300 selects the output unit with the highest activity level as the
10 most likely number desired by the user. System 300 then initiates the verification procedure (see FIG. 2), and, if successful, invokes dialer 360 to dial the number.

Network 500 thus does more than adjust the relative contribution of speech
recognition system 330 and model 320 in making each decision. It makes the
adjustments differentially for each number. This is desirable because name and
15 number combinations differ in the extent to which it helps for predictive model 320 to
override the decision of recognizer 330. When the user systematically and repeatedly
mispronounces a person's name or where the user correctly pronounces a person's
name but system 300 has matched the incoming speech to an orthographic model
because of improper pronunciation, system 300 will learn that predictive model 320
20 needs to be given more weight to adjust the model for these matching errors.

Integrator 350's neural network 500 is trained by the backward propagation learning algorithm, as is the case for the neural network 400 of the user's calling

behavior model 320 described previously. Networks 400 and 500 are trained separately, but network 500 uses a training set consisting of a set of telephone numbers and match quality pairs for input and a single telephone number for output.

When training neural network 500, the input units with a connection from calling

5 behavior model 400 for any of the telephone numbers having outputs from speech recognition system 330 are set to a fixed value, such as 0.8. This value corresponds to the maximum expected output from the calling behavior model network assuming a very likely call. All other units with inputs from calling behavior model 320 are set to 0.

10 When a name and telephone number is in the user's personal directory, but the user has never successfully dialed it by voice, there is no speech matching quality data available. In such cases a "dummy" training example is created that has a single telephone number and a speech match quality set to a fixed moderate value (e.g., 0.50), with the speech match quality set to 0 for all other numbers.

15 c. Call Processing for Training Neural Nets

FIGs. 6a and 6b show a flowchart of the procedure 600 used by the system 300 during incoming and outgoing telephone calls for recording information for training the neural networks 400 and 500. Once initiated, system 300 tests the type of call, whether incoming or outgoing (step 602). If the call is an incoming call, system 20 300 determines whether the call is not answered, answered by a voicemail system, or answered by a human (step 604). If the call was not answered the number of the calling party can be stored if the telephone has a caller ID system. If the call was not

answered but caller ID data is available, system 300 saves a record in the historical training database for training the calling behavior model (step 610), in this case the fact of the call being received (duration of call=0).

5 If the call was answered by a voicemail system (step 604), system 300 plays speech from a recorded answering message (step 616), and attempts to record a message. System 300 starts a timer and, when the message is complete, it determines the duration of the call (step 608), and saves a record in the historical database for the calling behavior neural network 400 (step 610). Control then passes to the beginning to wait for another call (step 602).

10 If the call was answered by a human (step 604), control passes to monitor the speech and attempt to recognize such phrases as "wrong number", "sorry, wrong", etc., indicating that the call is in error (step 612). Because of the tendency of speech recognition system 330 operating in this "word-spotting" mode to generate false alarms, the threshold for recognition of one of these phrases is set high, and the 15 phrase must also occur within a certain elapsed time (e.g., 20 seconds) after the beginning of the call.

20 In an alternative embodiment, the syntax of common dialogue interactions, such as "Can I speak to Debbie Heystek? No, there is no one here by that name" are encoded in the grammar of a speech recognition system. The system can also perform syntactic processing and, by assessing the likelihood of different interactions likely to indicate a wrong number, can detect a "wrong number" situation with increased accuracy.

If a wrong number is detected (step 612), control passes to the beginning of procedure 600 to wait for another call (step 602). If there is no wrong number detected, the duration of the call is determined (step 608), and a record is saved in the historical database for the calling behavior model (step 610).

5 If the call is an outgoing call (step 602), the system determines whether the number was dialed manually or by voice (step 622 in FIG. 6b). If the call is dialed manually (step 622), it is completed normally and system 300 determines the possibility of a "wrong number" in the manner discussed above (step 624). If the number is wrong (step 624), control passes to the beginning of procedure 600 (step 10 602 in FIG. 6a).

If system 300 does not detect a wrong number (step 624), it measures the duration of the call (step 626). When the call has been completed, system 300 stores a record of the call in the historical database for training the calling behavior neural network 400 (step 628). Control passes to the beginning of procedure 600 to wait for 15 another call (step 602 in FIG. 6a).

If the call is dialed by voice, speech recognition system 330 attempts to recognize the name (step 632). If successful, system 330 plays the name back to the user to verify (step 634). If the user's response to an attempt to verify is "Yes" (step 20 636), the call is placed (step 637), and monitored for "wrong number" indication (step 638). If so, control passes to the beginning of procedure 600 (step 602 in FIG. 6a).

If the number is not wrong, system 300 saves a record in the historical database to train the integrator neural network 500 (step 640).

When the call completes, the duration of the call is determined (step 642), and a record of the call is stored in the historical database for training the calling behavior neural network 400 (step 644). Control passes to the beginning of procedure 600 to wait for the next call (step 602 in FIG. 6a).

5 If the user's response to the verification is negative (step 636), the most recently entered record is deleted from the calling behavior model training database (step 646). The system then obtains the name with the next closest match (step 648), and verification continues (step 634).

10 Alternative procedures are also possible. For example, deleting the record (step 646) is not always necessary. Also, the user can respond with "Disconnect" instead of "Yes" or "No," with "Disconnect" causing the deletion of all historical records for the calling behavior network for the particular name and number. This response can be selected by the user when a particular number frequently overrides a desired number or numbers.

15 There is typically a limit to the number of records in the historical database for the neural network 400. To allow the historical database to keep storing records, old records must be purged. In the preferred embodiment, when at least five records exist for the same telephone number, the oldest of the records for that number is deleted. If not, the oldest record for the telephone number with the most records is deleted.

20 d. Calling Behavior Training Data Structure

FIG. 7 shows a data structure 700 of historical call information saved for use in training calling behavior neural network 400. The columns 710, 720, 730, 740,

750, 760, and 770 show data recorded for each call as a result of an incoming or outgoing call. A record 780 includes for each call:

- 1) a date of the call 710;
- 2) a day of the week 720 (0-6 records Monday through Sunday, respectively);
- 3) a time of day 730 (0 if midnight-6am, 1 if 6-9am, 2 if 9-12 am, 3 if 12-1, 4 if 1-4 pm, 5 if 4-6 pm, and 6 if 6-12 pm.);
- 4) a telephone number 740;
- 5) an indication 750 of whether the call was incoming or outgoing (0 if incoming, 1 if outgoing);
- 6) an indication 760 how the call was answered, if an incoming call 0 indicates not answered, 1 indicates answered by a voicemail system, and 2 indicates answered by a human; and
- 7) call duration 770.

15 Preferably, the duration of an answered call is measured in seconds. The duration of an unanswered call is measured by the number of rings. For an incoming call answered by a voicemail system or answering machine, the duration is the time elapsed from the end of the message played to the caller to the end of the message left by the caller. For all other calls the duration is measured from the beginning of the 20 connection to its end.

e. Integrator Training Data Structure

FIG. 8 shows a data structure 800 of historical call information saved for use

in training integrator neural network 500. Pairs of columns 820 and 825, 830 and 835, and 840 and 845, show three sets of numbers to be called and the corresponding quality of the match (similarity measure) between the incoming speech signal and the phonetic representation in the speech database managed by the speech recognition system 330. Column 850 shows the correct number that the user indicated in the verification procedure was correct by responding with "Yes." FIG. 8 shows only three pairs for clarity. An actual system would have 5-10 pairs saved per correct call.

f. Training Routine

FIG. 9 shows a flowchart of a procedure 900 used by voice-dialing system 300 to train neural networks 400 and 500. The procedure 900 is part of a system training controller 370 and is implemented in software.

When voice-dialing system 300 is installed on workstation 110, the user sets a parameter in system training controller 370 to indicate a daily time (e.g., 2 a.m.) that system training controller 370 uses for training both neural networks 400 and 500.

This time should be chosen to avoid periods when the computer is in use. When the appropriate time is reached, controller 370 tests to ensure that workstation 110 is not in use and is available for network training. If workstation 110 is in use, controller 370 waits until the recent history of microprocessor 140 usage is such that it is clear that adequate computation time is available without interfering with user activity.

When workstation 110 is available and training is initiated, controller 370 configures the network architecture for the calling behavior network (step 910). This is done by determining the number of names and associated telephone numbers in the

current directory, and constructing a network with the appropriate number of hidden units, output units, and connections between the input layer and hidden layer and between the hidden layer and output layer. The exact number of hidden units can be adjusted to yield the best generalization performance. Rules based on these 5 adjustments are encoded in the architecture configuration and learning part of the eventual product. The number of hidden units must be substantially less than the number of combinations of telephone numbers crossed with the alternative times, etc., so as to force the network to generalize.

Next, controller 370 trains the calling behavior network (step 920) and 10 configures the integrator network architecture 500 based on the number of names and associated numbers in the directory (step 930). Finally, the controller 370 trains the integrator network (step 940).

i. Training - Calling Behavior Network

FIG. 10 shows a flowchart describing the steps for training calling behavior 15 network 400 (see step 920 of FIG. 9). When initiated, controller 370 first builds a training set from historical training data shown in FIG. 7 (step 1010). This historical training data is stored in a database file on hard disk 160 (FIG. 1). Each record in the historical database is converted to an example for training. In addition, each number 20 in the personal directory is searched for in the historical database. If a record is not found, a training example is created for that number with a minimum selection probability.

The following fields are preferably defined in the set of training examples:

day, time, telephone number, selection probability, and the number of records for the given telephone number. The day and time are provided to the network input layer when training, and the telephone number is provided at the output layer of the network 400 for use by the learning algorithm. The selection probability is a 5 parameter that defines the probability that the example will, at any given cycle through the training procedure, be selected as a training example. It has a value between 0.002 to 1.00. The number of records, and thus examples for the given telephone number is needed to determine the extent to which the inputs are set to random values, rather than the actual day of the week and time of day. This is 10 necessary when relatively few records exist for a given telephone number, preventing the network from generalizing sufficiently to output the given telephone number if the input values are different from those for the existing records.

For example, if a single record exists of a call to a number at 12 noon on Tuesday, the network, if trained with only this data, would have a substantial output 15 value for the given telephone number only if the input specified noon on Tuesday. If 50 records existed for the same number at diverse times and days, a reasonable output value would be likely for that number with inputs at nearly any time and day. If 50 records existed for the same number, all at noon on Tuesday, the network would again be responsive for the given number only at that time and day, but quite legitimately 20 so, given the history of calls.

The selection probability is computed as follows:

$$P_{\text{selection}} = X_{\text{age}} * X_{\text{duration}},$$

where X_{age} ranges from 0.01 to 1.0 and $X_{duration}$ ranges from 0.20 to 1.00. $P_{selection}$ thus has possible values from 0.002 to 1.00.

X_{age} the number of days between the call being made and the network being trained. The values are assigned as follows: 0.01 if the call was made over a year ago, 0.02 if the call was made 181-365 days ago, 0.04 if the call was made 91-180 days ago, 0.08 if the call was made 31-90 days ago, 0.15 if the call was made 10-30 days ago, 0.30 if the call was made 4-9 days ago, 0.60 if the call was made 2-3 days ago, and 1.0 if the call was made yesterday.

The value of $X_{duration}$ depends on both the circumstances of the call and the actual duration. For outgoing calls or incoming calls answered by a human, $X_{duration}$ is assigned as follows: 1.0 if duration > 60 minutes; 0.8 if 11-60 minutes, 0.6 if 2-10 minutes, 0.4 if 30-119 seconds, and 0.20 if 15-30 seconds. If the call duration is less than 15 seconds, the record is discarded as unreliable.

For unanswered incoming calls, the call must ring at least twice or the record is discarded as unreliable. For unanswered incoming calls that are not answered that ring at least twice, when such a call is detected and the age of the call is three days or less, a search of following records is made to determine whether a later record exists of an outgoing call to the same number indicating that the call has been returned. For unanswered unreturned incoming calls with two or more rings that is three days old or less, in which the user has access to caller ID records that show who has called, $X_{duration}$ is as follows: 0.4 if 2-4 rings, 0.8 if 5-7 rings, and 1.0 if 8 or more rings. For unanswered incoming calls not meeting these conditions, $X_{duration}$ is as follows: 0.2 if

2-4 rings; 0.4 if 5-7 rings, and 0.6 if 8 or more rings.

For incoming calls that are answered by a voicemail system with a message left, the message must be at least five seconds long or the record is discarded as unreliable. For such calls a search of records is done to determine whether the call
5 has been returned, assuming that the message is five seconds or more in duration, and is no more than three days old. For unreturned calls meeting the criteria, $X_{duration}$ is as follows: 0.4 if the message is 5-15 seconds in duration, 0.8 if 16-60 seconds in duration, and 1.0 if 61 seconds or more in duration. For other calls, $X_{duration}$ is as follows: 0.2 if the message is 5-15 seconds in duration, 0.4 if 16-60 seconds, and 0.6
10 if 61 seconds or more.

After the training set has been constructed (step 1010), the set of weights for the connections between units of the network is then set to random values to initialize the network 400 (step 1020). A training example is then obtained from the training set (step 1030). This training example is the first in the set if the weights have just
15 been initialized. Otherwise, the next example in the set is selected. If there are no more examples in the set (as the result of the previous training example being the last in the set), the first example in the set is selected.

A calculation is then made to determine whether the example just selected is actually used to train the network on the current pass (step 1040). The selection
20 probability for the example is retrieved, and a random number from 0 to 1.0 is generated and compared with the selection probability. Only if the number is less than or equal to the selection probability is the example used.

For example, if the selection probability is 0.5, then the example is only used when the random number generated is from 0 to 0.5, or 50% of the time. If the selection probability is 0.1, then the example is only used when the random number is from 0 to 0.1, or 10% of the time. If the example is not used, control is passed to 5 obtain another training example (step 1030).

Otherwise, the network 400 is trained with the example and the accumulated error saved (step 1050). This is done by first providing the input of the network 400 with the appropriate input signals. These can be either the actual inputs for the example or, as suggested above, randomized inputs.

10 When training begins, a parameter known as the input randomization probability cutoff, P_{cutoff} , is calculated according to the following formula:

$$P_{cutoff} = N_{records} / N_{combinations},$$

where $N_{records}$ is the number of records for this number in the historical database, and $N_{combinations}$ is the number of input combinations, which equals the number of levels of 15 the day of week multiplied by the number of levels of the time of day input. For the network shown, $N_{combinations} = 7 \times 7 = 49$.

A random number from 0 to 1 is generated for each example and compared with the input randomization probability cutoff, P_{cutoff} . If the number is less than P_{cutoff} , a random number from 1 to 7 is generated to input to the time of day units, and 20 a separate random number from 1 to 7 is generated to input to the day of week units. If the random number is equal to or greater than P_{cutoff} , the actual inputs from the example are fed to the input units of the network.

Thus, for example, if only 1 record was available, P_{cutoff} is 1/49, or about 0.02, and the network 400 would be trained with a random date and time for 98% of the training trials (on average). For 49 available records, P_{cutoff} is $49/49=1.0$, and the network 400 would be trained with the actual date and time essentially all of the time.

5 Training is done by applying the example to the appropriate inputs and outputs of the network 400, then using the backward propagation learning algorithm to modify the values of the weights of the connections in the network 400. Details of the backward propagation algorithm are described in the Rumelhart, Hinton, and Williams paper, which was referred to above. In this training, a set of data is used
10 that includes both input and output data.

Thus, for example, a particular piece of data might consist of the day of the week and the time of day for inputs and a telephone number as output. The input data to the input layer are entered by setting the input unit matching the output from the example to 1.0, or "true," and setting all other input units to 0.0, or "false." Thus, in
15 the case of the day of the week "Tuesday", the input unit corresponding to "Tuesday" is set to 1.0, while the other 6 input units are set to 0.0.

The telephone number for each trial is then effectively applied to the output units using the following steps. First, information is applied to the inputs of the network and then allowed to propagate through the network to the output units. Next,
20 a calculation is made of the "error" of the network for each output unit by subtracting the actual output (activity level) of each unit from either 1.0, if the unit corresponds to the telephone number associated with the given trial, or 0.0. This error value is then

"propagated backward" through the earlier layers of the network 400, by systematically changing the values of the weights according to the backward propagation learning algorithm in such a manner as to reduce the error. A given set of data is applied repeatedly to a network 400 until overall error is reduced to the 5 point that the network 400 is considered trained.

The "accumulated error" is determined by summing the error for all output units across all training examples. The error for each unit is equal to the desired output value minus the actual output value. After training the network 400 with an example, a test is made of the result of the training thus far (step 1060). The 10 backward propagation learning algorithm is a "hill-climbing" algorithm. It uses a computation based on local information to seek a global minimum of error.

Such an algorithm can become "stuck," however. Networks may oscillate, continuing to learn for a short period but then falling back. The accumulated error after training is tested against a threshold level below which the network 400 is 15 considered fully trained. If the error is above the threshold and the number of training trials is below a maximum, the network 400 needs more training. If the error is above the threshold and the maximum allowed number of training trials have been reached, the network 400 is considered "stuck." In general, the complexity of the problem is low and it is unlikely that the network 400 will become stuck. Because certain sets of 20 random weight values can cause a network to become stuck even with problems of low complexity, it is necessary to test for this condition and respond to it.

If network 400 needs more training (step 1060), control returns to obtain

another training example (step 1030). If the network 400 is "stuck," control passes to initialize the weights and begin the training process from the beginning (step 1020).

If the network 400 has its accumulated error below the threshold, then the training is completed.

5

ii. Training - Integrator Network

FIG. 11 is a flowchart describing the procedure for training the integrator network (see step 940 in FIG. 9). When initiated, the training controller 370 first builds a training set from the historical training data (step 1110). Each record in the historical database is used as an example for training, with the data read into temporary RAM 150 to allow rapid training. In addition, a list of all of the telephone numbers referred to in the historical database is created, and its contents matched against all numbers in the personal directory to determine those numbers entered into the directory for which there is no historical database record. A training example for each of these numbers is also created in RAM 150, with the match quality set to 0.8 for the number in question and 0 for other numbers.

15 The set of weights for the connections between units of network 500 is then set to random values to initialize network 500 (step 1120). A training example is then obtained from the training set (step 1130). This training example is the first in the set if the weights have just been initialized. Otherwise, the next example in the set is selected. If there are no more examples in the set (as a result of the previous training example being the last in the set), the first example in the set is selected.

20 Network 500 is then trained with the example and the accumulated error saved

(step 1140). This is done by first providing those input units of the network for which there exists a speech match quality score. The input values to all other input units is set to 0.

5 Training is done by applying the example to the appropriate inputs and outputs of network 500, then using the backward propagation learning algorithm to modify the values of the weights of the connections in the network.

The "accumulated error" is determined by summing the error for all output units across all training examples. Error for each unit is equal to the desired output value minus the actual output value.

10 After training network 500 with an example, a test is made of the result of the training by comparing the accumulated error after training against a threshold level below which network 500 is considered fully trained (step 1150). If the error is above the threshold and the number of training trials is below a maximum, network 500 needs more training. If the error is above the threshold and the maximum allowed 15 number of training trials have been reached, network 500 is considered "stuck."

If network 500 needs more training (step 1150), control is passed to obtain another training example (step 1130). If network 500 is "stuck," control is passed to re-initialize the weights and begin the training process from the beginning (step 1120). If network 500 has its accumulated error below the threshold, then the training 20 is completed.

g. Procedure for Modifying Personal Directory

FIG. 12 shows a flowchart of the preferred procedure 1200 to control the

process used when the user modifies the directory of names and associated telephone numbers. The voice-dialing system 300 includes a component having software corresponding to procedure 1200. Microprocessor 140 executes procedure 1200 to modify the personal directory.

5 First, the type of modification is determined (step 1205). The user can add a new record, delete an existing record, or modify an existing record. If the modification requires adding a new record, the name and associated telephone number are entered by the user using keyboard 190 and a graphical interface on display 180, and mouse 195. The name is then added to the list of names contained in
10 the speech recognition system 330 (step 1210). The name, and appropriate codes for connecting the name and telephone number with the software that provides an interface with neural networks 400 and 500, are then added to a temporary store for use until the networks 400 and 500 are trained to make use of the new name directly (step 1215). The procedure is then finished.

15 If the modification involves deleting a record (step 1205), the name is removed from the list of names contained in the speech recognition system 330 (step 1220). System 300 then searches the databases containing historical data for training calling behavior network 400 and integrator network 500, and deletes all records that refer to the number being deleted (step 1225). The connections to the inputs of
20 networks 400 and 500 that refer to that number are then disconnected (step 1230), so that there will be no activity in networks 400 and 500 for that number. The procedure is then finished.

If the modification involves modifying an existing record (step 1205), a test is made of whether the modification is to a name or a number (step 1240). If the modification is to a name, the name in the speech recognition system 330 is modified as appropriate (step 1245). The procedure is then finished.

5 If the modification is to a number, system 300 then searches the databases containing historical data for training both calling behavior network 400 and integrator network 500, and deletes all records that refer to the old number being deleted (step 1250). The connections to the inputs of networks 400 and 500 that refer to that old number are then disconnected (step 1255), so that there will be no activity
10 in the networks 400 and 500 for that number. The new number is added to a temporary store for use until networks 400 and 500 are trained to make use of the new number directly (step 1260). The procedure is then finished.

B. PBX System

1. System Architecture

15 FIG. 13 shows an alternate embodiment in which the voice-dialing system according to the present invention is implemented as a server 1300 for a PBX system 1310 to provide voice-dialing service for all of the telephone users at a particular site. Server 1300 is connected via a high-speed digital connection to PBX system 1310 that contains a number of telephone lines connected to telephones 1320a-d. A typical
20 PBX would have tens to hundreds of these lines. PBX system 1310 may be a Northern Telecom Meridian 1 PBX system, with a T1 digital connection between server 1300 and PBX system 1310. Server 1300 consists primarily of the same

hardware components 140-195 illustrated in FIG. 1.

In operation, server 1300 maintains, in the form of stored weights on hard disk 160, a separate neural network architecture and memory of calling behavior history for each user (telephones 1320a-d). Video display 180, keyboard 190, and mouse 195 are for maintenance of server 1300 and could be dispensed with, particularly if server 1300 was connected to a local area network such that maintenance of server 1300 could be done using a remote workstation over the local area network. PBX system 1310 is also connected to the public switched telephone network.

Floppy disk 170 is for loading software and could be dispensed with if software were loaded over a local area network.

When a user picks up a handset of one of telephones 1320a-d, weights of the calling behavior neural network for that telephone are read into RAM 150. The calling behavior neural network for that telephone is then executed for the given time and date as described above.

15 2. Software Components

FIG. 14 shows a block diagram of the software components for the PBX-based voice-dialing system 1400 implemented using server 1300. Voice-dialing system 1400 consists of four primary components: (1) component 1410 that models the user's calling behavior based on the user's personal history of calls, (2) component 1420 models the calling behavior of groups of people at a physical or virtual site based on a sitewide history of calls between numbers that are defined in particular categories, (3) speech recognition system 1430, which may use conventional

techniques, a neural network, or a hybrid approach, and (4) integrator component 1440 that integrates the outputs of the first three components to produce the best estimate of the name (and number) the user desires to call.

Also included is a telephone dialer 1450 that looks up the actual telephone number in a table and dials it. Finally, a system training controller 1460 trains the networks 1410, 1420, and 1440. Both calling behavior model component 1410 and the integrator 1440 are preferably implemented as neural networks and have historical training data 1405 and 1435, respectively, that are maintained to continue training the neural networks when appropriate.

The personal history calling behavior model 1410 is identical to the component 320 used in voice-dialing 300 (FIG. 3), and its architecture is shown in FIG. 4. The category-based call behavior model 1420 also tries to predict the likelihood that a given number will be called for a given calling number, but the method used is very different.

In general, all telephone numbers in the PBX system at a given site are divided into categories according to the organization of the institution. For example, each department or other group in an institution may be a different category. Processor 1400 records all calls from one PBX number to another over a period of time and stores them in a database of historical training data 1415, identifies the appropriate category for each incoming and outgoing number, and then trains the network. The network can therefore give an input category, generate an output signal for each category that predicts the likelihood of a number in the category being called.

The output category is then converted to specific numbers and provided as input to integrator 1440.

The speech input from the person saying the name of the number to be dialed is processed by speech processor component 1430 and then fed into integrator network 1440. The historical training data 1435 for integrator network 1435 5 preferably has the same as the fields for speech recognition match quality as shown in FIG. 8.

These values are used together with the output values from the category-based calling behavior model component 1420 to train integrator network 1440. A fixed 10 number, such as 0.9, which represents the assumed value of the output of personal history calling behavior network 1410 just after a call is made, is also used to train integrator network 1440.

3. Category-Based Calling Behavior Neural Network

FIG. 15 shows the architecture of category based calling behavior neural network 1500. Neural network 1500 is used by component 1420 in the block diagram 15 of FIG. 14. Network 1500 is a three-layer feedforward neural network consisting of an input layer 1510, hidden layer 1520, and output layer 1530. The network architecture is similar to network 400 shown in FIG. 4, except for the differences indicated below.

20 Input layer 1510 consists of three groups of input units. One group 1512 encodes the current day of the week, and consists of seven units, one for each day of the week. The second group 1514 encodes the current time of day and consists of

seven units (midnight-6am, 6-9am, 9-12 am, 12-1, 1-4 pm, 4-6 pm, 6-12 pm). The third group 1516 of input units consists of one unit for each organizational category defined in the voice-dialing system 1400. Training is done from a set of examples of calls in which the telephone numbers for the calling and called numbers are converted to categories. Thus, each training example is composed of a day of the week, time of day, calling number category, and called number category.

5 An alternative embodiment uses an architecture, and resulting training data, that eliminates the day of the week and time of day inputs to reduce computational requirements. Once the network 1500 is trained, it is activated by providing it with the current day of week and time of day, and the organizational category of the 10 number that the particular call is from. Network 1500 predicts the likelihood of calls to numbers in particular categories, and, by means of table lookup, the likelihood of calls to specific numbers.

15 An alternative method of implementation, the model is constructed by creating a table that maintains a count of calls in each $N \times N$ combination, where N is the number of categories of numbers. A given call from a number in one category to another results in an increment in the appropriate count. When a call is initiated from a particular number, the category is determined by a table lookup, and a number indicating the relative likelihood that the call would be made to a number in particular 20 categories is completed from the normalized counts and provided as outputs from the model.

Training is done in a manner similar to that for integrator 350 for voice-

dialing system 300 (see FIG. 5), with two exceptions. First, two predictive models are used instead of one, and second, integrator 1440 is used to weight the contribution of the two predictive models on the basis of which is a better predictor in a given case, and weighting the contribution of the speech processor 1430. The integrator 5 network could also be replaced with a simple numerical algorithm that contained fixed weights, that simply computed a weighted average of the outputs of the three components, albeit with lesser performance.

In an alternative embodiment, a model of the frequency of calls to particular numbers at the given site is used in addition to, or instead of, the other models. Such 10 a model is implemented with a neural network architecture that is similar to calling behavior model network 400 shown in FIG. 4, with the output units representing the entire set of numbers at the site rather than the numbers in a personal directory. This network has its historical training data created in a manner analogous to calling behavior model shown in FIGS. 6a and 6b.

15 Training the network is also done in a manner analogous to the calling behavior model as shown in FIG. 10, except that the steps of selecting an example probabilistically according to the number of days since the example calls are made are not used. All examples, however, are used.

An alternative method to implement the model is to simply count the number 20 of calls to particular destination numbers and produce an output reflected by this proportion of calls. This would not capture interactions of calling patterns with the day of the week and the time of day, however.

Many alternatives exist for the overall architecture of the system, given the three possible models of personal history calling behavior model, category-based calling behavior model, and frequency-of-destination number model. An architecture can also be constructed by combining the speech recognition component 1430 with 5 any one of the other calling behavior models, any two of them, or all three. The category-based calling behavior model component 1420 and frequency-of-destination number model require a relatively short period of time for training. Once trained, these networks are available for all users to train the integrator network for specific users and to predict the call likelihoods for specific users. In contrast, each user has a 10 personal history calling behavior model and integrator network devoted specifically to that user.

FIG. 16 shows a block diagram of an alternative architecture for the voice-dialing system that was previously shown in FIG. 3. FIG. 16 shows software system 1600 executed by microprocessor 140. Software system 1600 may be stored on hard 15 disk 160. This alternative architecture is typically computationally faster than that shown in FIG. 3.

System 1600 consists of two principal components: a model of the user's calling behavior 1620 and a speech recognition system 1630. System 1600 also includes a telephone dialer 1640 that looks up the actual telephone number in a table 20 and dials it. System training controller 1650 trains the calling behavior model 1620, using historical training data 1610. System training controller 1650 works as described in FIGs. 9 and 10, except that only the calling behavior network, not the

integrator, is trained. Calling behavior model 1620 preferably includes a neural network and uses historical training data 1610 that is maintained to continue training the neural network when appropriate.

When a user picks up the handset of telephone 120 or dials in to the 5 workstation 110 from a remote telephone and identifies himself or herself, microprocessor 140 reads in the weights of the user's calling behavior model 1620 from hard disk 160 and determines the current time and day of the week and provides this to the calling behavior model. The calling behavior model is then computed from these inputs and provides at its outputs a prediction of the likelihood that the user will 10 call each telephone number included in the model. These predictions are provided to the speech recognition system 1630.

When the user speaks the name of the person to be called, speech recognition system 1630 processes the input speech data and attempts to match it against the set of stored representations, typically sequences of phonemes, that represent each name 15 in the database. Unlike the case of the system 300, however, speech recognition system 1630 also takes into account context information, in particular the likelihood of calling particular telephone numbers, that is associated with each name in the database. The provision of this information to the speech recognition system 1630 allows that system to quickly eliminate processing for alternatives that have both a 20 low likelihood of being called and have low similarity to the target sequence of phonemes, resulting in increased computational efficiency and also faster response time.

This alternative architecture requires the use of a speech recognition system, which can be a commercially available speech recognition "engine," that includes the capability of having word and phrase recognition driven by higher-level context information. An example of such a system is the "Watson" speech recognition system 5 that is sold as a product by AT&T in the United States.

Speech recognition system 1630 produces sets of floating point numbers, each representing the extent to which there is a match between the speech input and the stored representation for the name associated with each telephone number, taking into account the biasing information provided by the calling behavior model.

10 Commercially available speech recognition engines such as the "Watson" engine referred to above typically produce an output consisting of the "N best" matches of names in the database for which the match (including biasing information) was above a given threshold value, with a quality measure for each. The number with the highest quality can be dialed immediately, or the list, in order of quality, can be used 15 for selecting the best name and number to be provided to the user for confirmation in the protocol indicated in FIG. 2.

The alternative architecture shown in FIG. 16 can also be applied to the PBX server embodiment shown in FIGs. 13-15. In this case the block diagram of FIG. 14 is modified in a manner similar to the modification of FIG. 3, with both the personal 20 history calling behavior model 1410 and the category based calling behavior model 1420 providing its outputs to a speech recognition system 1430. Speech recognition system 1430 would require a speech engine capable of being context driven, and its

output would feed directly into telephone dialer 1450, with integrator 1440 and historical training data 1435 being eliminated.

C. Conclusion

The present invention thus provides a faster and more accurate voice-dialing system by building and maintaining a model of the calling behavior of an individual and using this model to increase the performance of the automatic speech recognition system that matches incoming spoken names with names stored in a directory. The voice-dialing system includes a component that models the user's calling behavior, a component that processes incoming speech and matches it against representations of the names in the directory, and a component that integrates the outputs of the first two components to produce the name that the user most likely desires to call. The user calling behavior model component consists of a multilayer feedforward neural network that uses the backward propagation learning algorithm. The inputs to the neural network accept the current date and time, while the output of the network provides a signal for each telephone number in the directory. The neural network is trained with a database of telephone numbers that have been received or called by the user along with the date and time of the call, whether the call was incoming or outgoing, how it was answered, and the duration of the call. Full retraining of the network is preferably done daily, typically during the early morning hours or when the network is not in use. Example telephone calls are selected for training probabilistically, with the probability that a given call in the set will be used in a given training trial a monotonically decreasing function of the time since that call was

made.

The component of the system that integrates the outputs of the first two components also consists of a multilayer feedforward neural network using backward propagation. The inputs to this neural network include one input for each telephone number in the directory from the output of the calling behavior model network, and one input for each telephone number from the output of the speech recognizer. The system is trained by a database of telephone numbers that have been dialed by voice, with each training example including those names (and associated numbers) that most closely match the speech that resulted in the call.

The present invention also facilitates fast and accurate voice-dialing within a site using a PBX system. According to this approach, the voice-dialing system uses three neural networks for a given individual. One neural network is common to all individuals of the organization and implements a predictive model of calling between individuals of the organization. This neural network is a multilayer feedforward neural network that uses the backward propagation learning algorithm. Every telephone number in the organization is associated with a category, with the category assignment made according to the structure of the organization at the site. The common network contains an input unit for each category and an output unit for each category. Training of the network is done with a list of telephone calls from one number to another over a relatively short period of time, such as a week, with each number converted to the appropriate category for that number before being applied to train the network.

The second neural network creates a model of the calling behavior for the specific individual, and its architecture and method of training is similar to that for user's calling behavior model, except that examples are not selected probabilistically based on the elapsed time since the example call was made. The third neural network 5 integrates together information from the first two networks and the speech recognition system to predict the likelihood of calls to particular numbers.

The foregoing description of a preferred embodiment of the invention has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Modifications and 10 variations are possible in light of the above teachings or may be acquired from practice of the invention. The scope of the invention is defined by the claims and their equivalents.

WHAT IS CLAIMED IS:

1. A method for assisting voice-dialing comprising the steps of:
receiving voice input from a user representing a name corresponding to a
desired telephone number;
selecting stored names that most closely match the voice input;
predicting a likelihood of the user calling telephone numbers based on a
model of the user's calling behavior; and
determining the desired telephone number according to the predicted
likelihood of the user calling the telephone number corresponding to each selected
name.

2. The method of claim 1, wherein the model of the user's calling
behavior includes weights determined from previous calls by the user to at least one
of the telephone numbers, and wherein the determining step includes the substep of:
applying the weights to order telephone numbers corresponding to the selected
names.

3. The method of claim 1, wherein the model of the user's calling
behavior includes weights determined from previous calls that the user received from
at least one of the telephone numbers, and wherein the determining step includes the
substep of:
applying the weighting factors to order telephone numbers corresponding to

the selected names.

4. The method of claim 1, wherein the determining step includes the substep of:

generating a set of the telephone numbers that are most likely the desired telephone number.

5. The method of claim 1, wherein the determining step includes the substep of:

ordering a set of the telephone numbers according to the predicted likelihood that each telephone number in the set is the desired telephone number.

6. The method of claim 2, wherein the determining step includes the substep of:

ordering the telephone numbers associated with the selected names according to the predicted likelihood that each telephone number in the set is the desired telephone number.

7. The method of claim 3, wherein the determining step includes the substep of:

ordering the telephone numbers associated with the selected names according to the predicted likelihood that each telephone number in the set is the desired

telephone number.

8. The method of claim 4, wherein the generating step includes the substeps of:

prompting the user to select one of the telephone numbers from the set; and
initiating a telephone call to the selected telephone number.

9. The method of claim 1 further comprising the step of:

dialing the desired telephone number.

10. The method of claim 1 further comprising the step of:

outputting the desired telephone number in a manner perceptible to the user.

11. The method of claim 1, wherein the model of the user's calling behavior comprises an abstract representation based on the user's environment and actions with respect to initiating telephone calls, and wherein the predicting step includes the substep of:

examining the abstract representation for indications that the user intends to call each of the telephone numbers.

12. The method of claim 1, wherein the model of the user's calling

behavior comprises an adaptive model that is alterable based on the user's environment and actions with respect to initiating telephone calls, and wherein the predicting step includes the substep of:

examining the adaptive model for indications that the user intends to call each of the telephone numbers.

13. The method of claim 1, wherein the model of the user's calling behavior comprises a neural network and wherein the predicting step includes the substep of:

examining the neural network for indications that the user intends to call each of the telephone numbers.

14. The method of claim 1, wherein a speech recognition system is used to receive the voice input, and wherein the determining step includes the substep of:

integrating the calling behavior model with a second model of accuracy measures for the speech recognition system.

15. The method of claim 14, wherein the second model includes integration factors determined from previous calls by the user to at least one of the telephone numbers, and wherein the integrating step includes the substep of:

applying the integration factors to select the desired telephone number.

16. The method of claim 15, wherein the applying step includes the substeps of:

generating a set of the telephone numbers associated with the selected names;
and

ordering the set of the telephone numbers in accordance with the integration factors with telephone numbers determined most likely to be the desired telephone number ahead of telephone numbers determined less likely to be the desired telephone number.

17. The method of claim 14, wherein the second model comprises an abstract representation based on the user's environment and actions with respect to initiating telephone calls, and wherein the integrating step includes the substep of:

examining the abstract representation for indications that the user intends to call each of the telephone numbers.

18. The method of claim 14, wherein the second model comprises an adaptive model that is alterable based on the user's environment and actions with respect to initiating telephone calls, and wherein the integrating step includes the substep of:

examining the adaptive model for indications that the user intends to call each of the telephone numbers.

19. The method of claim 14, wherein the second model of integration data is a neural network, and wherein the integrating step includes the substep of:

examining the neural network for indications that the user intends to call each of the telephone numbers.

20. The method of claim 14, wherein the second model is a fixed procedure, and wherein the integrating step includes the substep of:

combining the telephone numbers corresponding to the selected names with the predictions of the likelihood of the user calling each of the telephone numbers associated with the selected names.

21. The method of claim 1, further comprising the step of:
training the model of the user's calling behavior with previous calls from the user to each of the telephone numbers.

22. The method of claim 1, further comprising the step of:
training the model of the user's calling behavior with previous calls received by the user from each of the telephone numbers.

23. The method of claim A14, further comprising the step of:
training the second model with accuracy measures for the speech recognition system corresponding to previous calls.

24. The method of claim 4, wherein the generating step includes the substeps of:
- selecting a name associated with one of the telephone numbers in the set;
- presenting the selected name to the user; and
- waiting for a response from the user indicating whether the selected name corresponds to the desired telephone number.
25. The method of claim 24, wherein the waiting step includes the substep of:
- determining whether a predetermined period of time has passed since the user was presented with the selected name.
26. The method of claim 24, wherein the waiting step includes the substep of:
- interpreting a lack of response as meaning that the selected name corresponds to the desired telephone number.
27. The method of claim 1, further comprising the step of:
- building a training set including information related to at least one previous call.
28. The method of claim 27, further comprising the step of:

at a predetermined time, modifying the model of the user's calling behavior in accordance with the training set.

29. The method of claim 28, wherein the model of the user's calling behavior includes weights determined from previous calls by the user to at least one of the telephone numbers, and wherein the modifying step includes the substep of:
altering the weights of the user's calling behavior model to reflect the information related to the previous call.

30. The method of claim 14, wherein a speech recognition system is used to receive the voice input, further comprising the step of:
building a training set including information on at least one previous call and an accuracy measure of the speech recognition system in selecting a stored name corresponding to the voice input that corresponds to the previous call.

31. The method of claim 30, further comprising the step of:
at a predetermined time, modifying the second model in accordance with the training set.

32. The method of claim 31, wherein the modifying step includes that substep of:
altering the second model in accordance with the accuracy measure of the

previous call in the training set.

33. A system for assisting voice-dialing comprising:

means for receiving voice input from a user representing a name corresponding to a desired telephone number;

means for selecting stored names that most closely match the voice input;

means for predicting a likelihood of the user calling telephone numbers based on a model of the user's calling behavior; and

means for determining the desired telephone number according to the predicted likelihood of the user calling the telephone number corresponding to each selected name.

34. The system of claim 33, wherein the model of the user's calling behavior includes weights determined from previous calls by the user to at least one of the telephone numbers, and wherein the determining means includes:

means for applying the weights to order telephone numbers corresponding to the selected names.

35. The system of claim 33, wherein the model of the user's calling behavior includes weights determined from previous calls that the user received from at least one of the telephone numbers, and wherein the determining means includes:

means for applying the weighting factors to order telephone numbers corresponding to the selected names.

36. The system of claim 33, wherein the determining means includes:
means for generating a set of the telephone numbers that are most likely the
desired telephone number.

37. The system of claim 33, wherein the determining means includes:
means for ordering a set of the telephone numbers according to the predicted
likelihood that each telephone number in the set is the desired telephone number.

38. The system of claim 34, wherein the determining means includes:
means for ordering the telephone numbers associated with the selected names
according to the predicted likelihood that each telephone number in the set is the
desired telephone number.

39. The system of claim 35, wherein the determining means includes:
means for ordering the telephone numbers associated with the selected names
according to the predicted likelihood that each telephone number in the set is the
desired telephone number.

40. The system of claim 36, wherein the generating means includes:
means for prompting the user to select one of the telephone numbers from the
set; and
means for initiating a telephone call to the selected telephone number.

41. The system of claim 33 further comprising:

means for dialing the desired telephone number.

42. The system of claim 33 further comprising:

means for outputting the desired telephone number in a manner perceptible to the user.

43. The system of claim 33, wherein the model of the user's calling behavior comprises an abstract representation based on the user's environment and actions with respect to initiating telephone calls, and wherein the predicting means includes:

means for examining the abstract representation for indications that the user intends to call each of the telephone numbers.

44. The system of claim 33, wherein the model of the user's calling behavior comprises an adaptive model that is alterable based on the user's environment and actions with respect to initiating telephone calls, and wherein the predicting means includes:

means for examining the adaptive model for indications that the user intends to call each of the telephone numbers.

45. The system of claim 33, wherein the model of the user's calling

behavior comprises a neural network and wherein the predicting means includes:

means for examining the neural network for indications that the user intends to call each of the telephone numbers.

46. The system of claim 33, wherein the receiving means includes:

a speech recognition system,

and wherein the determining means includes:

means for integrating the calling behavior model with a second model of accuracy measures for the speech recognition system.

47. The system of claim 46, wherein the second model includes integration factors determined from previous calls by the user to at least one of the telephone numbers, and wherein the integrating means includes:

means for applying the integration factors to select the desired telephone number.

48. The system of claim 47, wherein the applying means includes:

means for generating a set of the telephone numbers associated with the selected names; and

means for ordering the set of the telephone numbers in accordance with the integration factors with telephone numbers determined most likely to be the desired telephone number ahead of telephone numbers determined less likely to be the desired

telephone number.

49. The system of claim 46, wherein the second model comprises an abstract representation based on the user's environment and actions with respect to initiating telephone calls, and wherein the integrating means includes:

means for examining the abstract representation for indications that the user intends to call each of the telephone numbers.

50. The system of claim 46, wherein the second model comprises an adaptive model that is alterable based on the user's environment and actions with respect to initiating telephone calls, and wherein the integrating means includes:

means for examining the adaptive model for indications that the user intends to call each of the telephone numbers.

51. The system of claim 46, wherein the second model of integration data is a neural network, and wherein the integrating means includes:

means for examining the neural network for indications that the user intends to call each of the telephone numbers.

52. The system of claim 46, wherein the second model is a fixed procedure, and wherein the integrating means includes:

means for combining the telephone numbers corresponding to the selected

names with the predictions of the likelihood of the user calling each of the telephone numbers associated with the selected names.

53. The system of claim 33, further comprising:
means for training the model of the user's calling behavior with information on previous calls to or from each of the telephone numbers.

54. The system of claim 46, further comprising:
means for training the second model with accuracy measures for the speech recognition system corresponding to previous calls.

55. The system of claim 36, wherein the generating means includes:
means for selecting a name associated with one of the telephone numbers in the set;
means for presenting the selected name to the user; and
means for waiting for a response from the user indicating whether the selected name corresponds to the desired telephone number.

56. The system of claim 55, wherein the waiting means includes:
means for determining whether a predetermined period of time has passed since the user was presented with the selected name.

57. The system of claim 55, wherein the waiting means includes:

means for interpreting a lack of response as meaning that the selected name corresponds to the desired telephone number.

58. The system of claim 33, further comprising:

means for building a training set including information related to at least one previous call.

59. The system of claim 58, further comprising:

means for modifying, at a predetermined time, the model of the user's calling behavior in accordance with the training set.

60. The system of claim 59, wherein the model of the user's calling behavior includes weights determined from previous calls by the user to at least one of the telephone numbers, and wherein the modifying means includes:

means for altering the weights of the user's calling behavior model to reflect the information related to the previous call.

61. The system of claim 46, wherein the receiving means includes:

a speech recognition system.

62. The system of claim 61, further comprising:

means for building a training set including information on at least one previous call and an accuracy measure of the speech recognition system in selecting a stored name corresponding to the voice input that corresponds to the previous call.

63. The system of claim 62, further comprising:

means for modifying, at a predetermined time, the second model in accordance with the training set.

64. The system of claim 63, wherein the modifying means includes:

means for altering the second model in accordance with the accuracy measure of the previous call in the training set.

65. A method for initiating telephone calls by voice, comprising the steps of:

activating, in response to a determination that a user intends to initiate a new telephone call, a calling behavior model to predict a likelihood of the user intending to call each one of a predetermined set of telephone numbers;

receiving a voice input including a sequence of sounds that represent a name spoken by the user and corresponding to a telephone number that the user desires to call;

selecting names from a directory including voice data representing names associated with the predetermined set of telephone numbers, based on a comparison of the sequence of sounds from the voice input and voice data of the directory; and

integrating the selection of telephone numbers corresponding to the selected names from the directory with the predictions of the likelihood of the user calling each of the telephone numbers associated with the selected names from the directory to identify one of the telephone numbers most likely to be the desired telephone number.

66. The method of claim 65, wherein the activating step includes the substep of:

stimulating a category-based calling behavior model to predict a likelihood of the user intending to call each one of a predetermined set of institutional directory telephone numbers, and

wherein the integrating step includes the substep of:

joining the predictions of the likelihood of the user intending to call each one of the institutional directory telephone numbers with the predictions of the likelihood of the user calling each of the telephone numbers associated with the selected names from the directory and with the telephone numbers associated with the selected names from the directory.

67. The method of claim 66, wherein the category-based calling behavior model comprises an abstract representation based on the user's environment and actions with respect to initiating telephone calls, and wherein the joining step includes:

examining the abstract representation for indications that the user intends to call each of the telephone numbers.

68. The method of claim 66, wherein the category-based calling behavior model comprises an adaptive model alterable in response to the user's environment and actions with respect to initiating telephone calls, and wherein the joining step includes the substep of:

examining the adaptive model for indications that the user intends to call each of the telephone numbers.

69. The method of claim 66, wherein the category-based calling behavior

model comprises a neural network, and wherein the joining step includes the substep of:

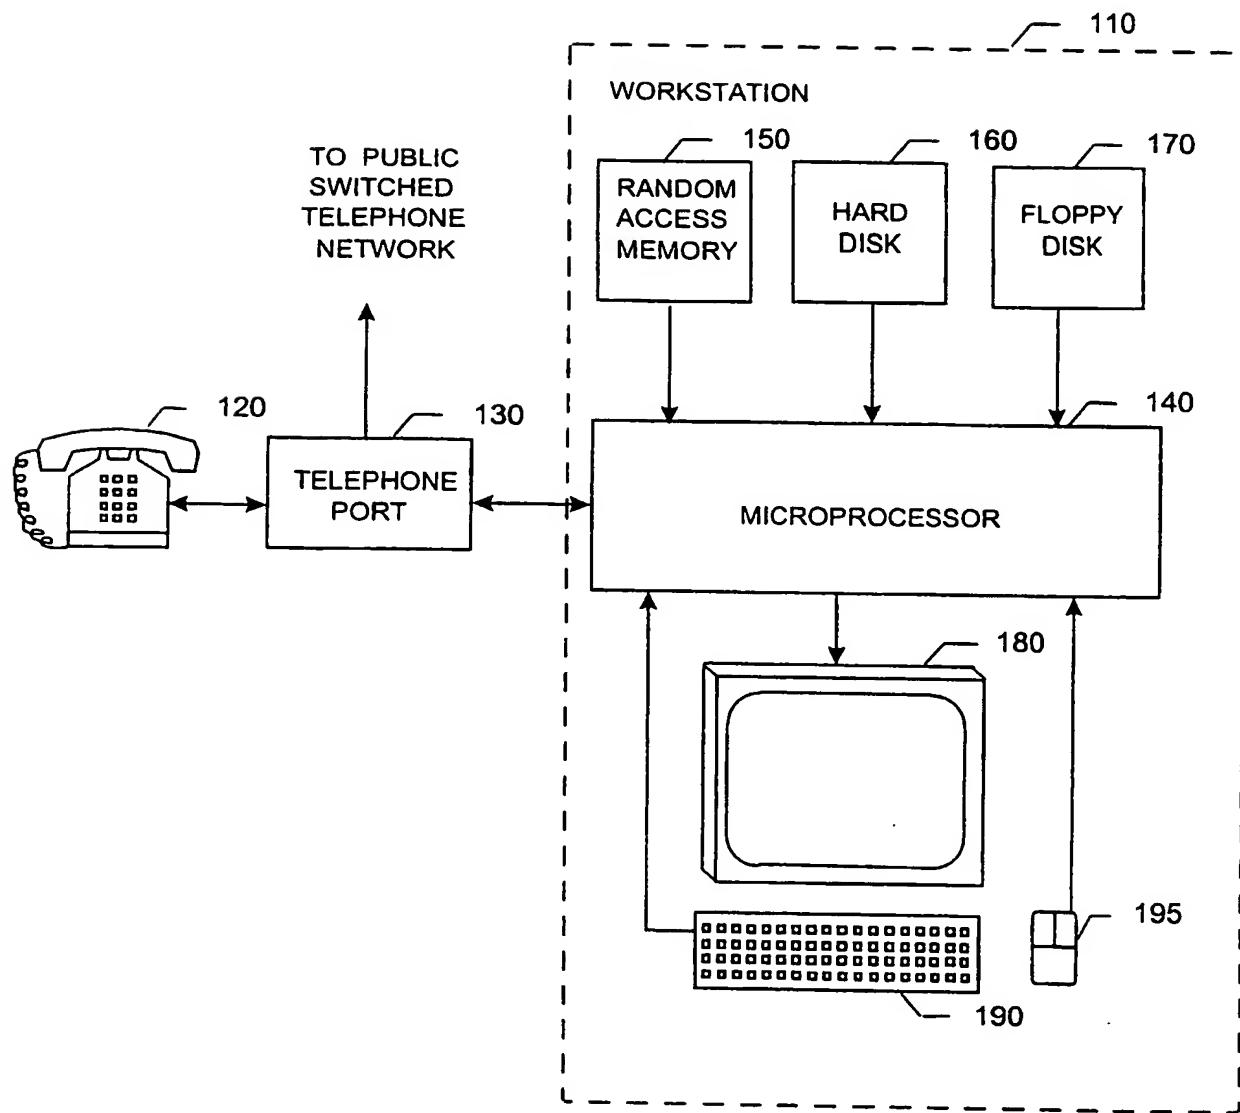
examining the neural network for indications that the user intends to call each of the telephone numbers.

70. The method of claim 66, further comprising the step of:
training the category-based calling behavior model with historical calling data based on previous times of calls among institutional directory telephone numbers .

71. A method for assisting voice-dialing comprising the steps of:
receiving voice input from a user representing a name corresponding to a desired telephone number;
predicting a likelihood of the user calling telephone numbers based on a model of the user's calling behavior; and
determining the desired telephone number according to the predicted likelihood of the user calling the telephone number corresponding to stored names that most closely match the voice input.

72. Voice-dialing apparatus, comprising:
a receiver configured to receive voice input from a user representing a name corresponding to a desired telephone number;
a predicting component configured to predict a likelihood of the user calling

telephone numbers based on a model of the user's calling behavior; and
a determining component configured to determine the desired telephone
number according to the predicted likelihood of the user calling the telephone number
corresponding to stored names that most closely match the voice input.

100**FIG. 1****SUBSTITUTE SHEET (RULE 26)**

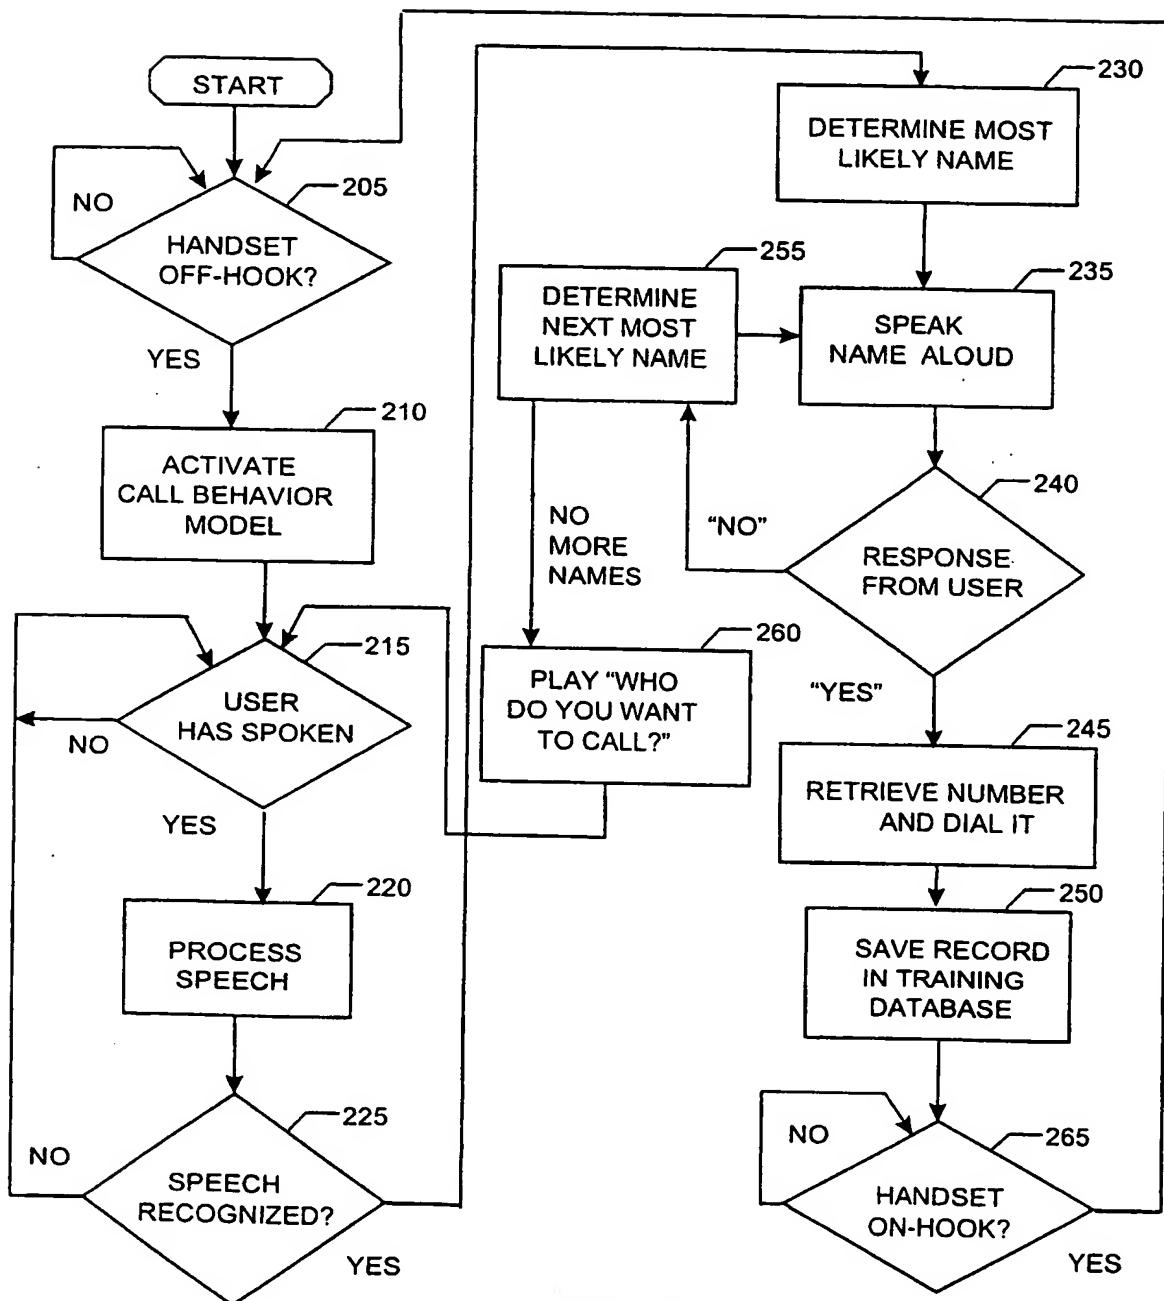
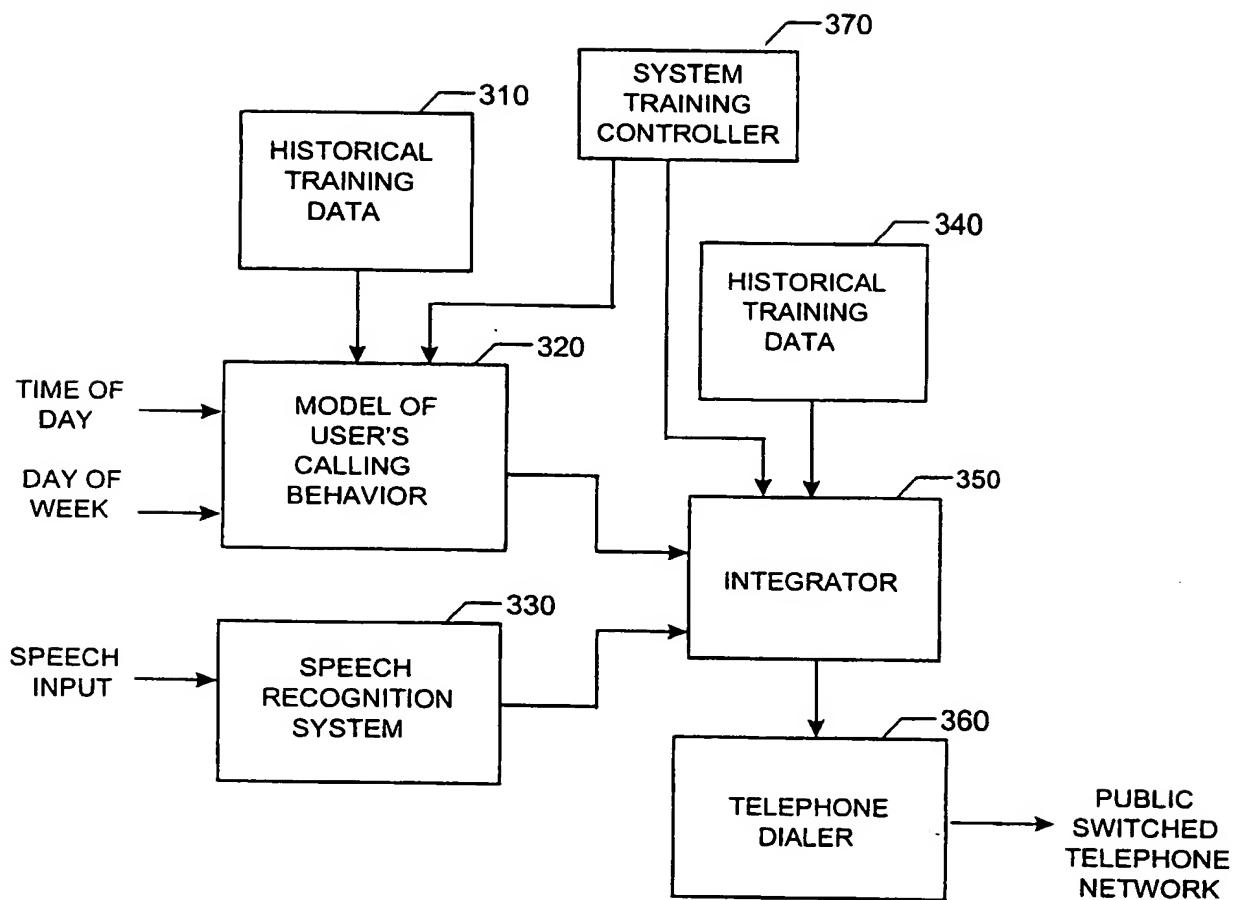
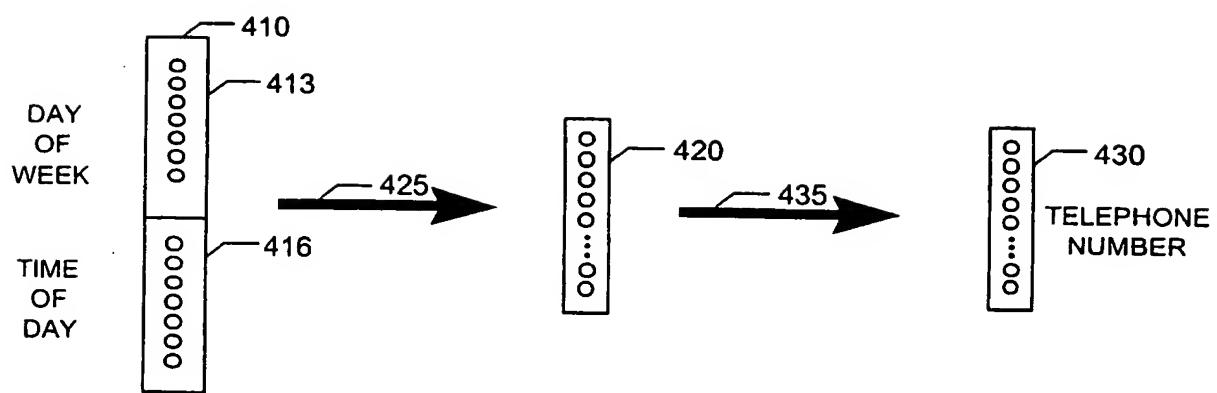
200

FIG. 2

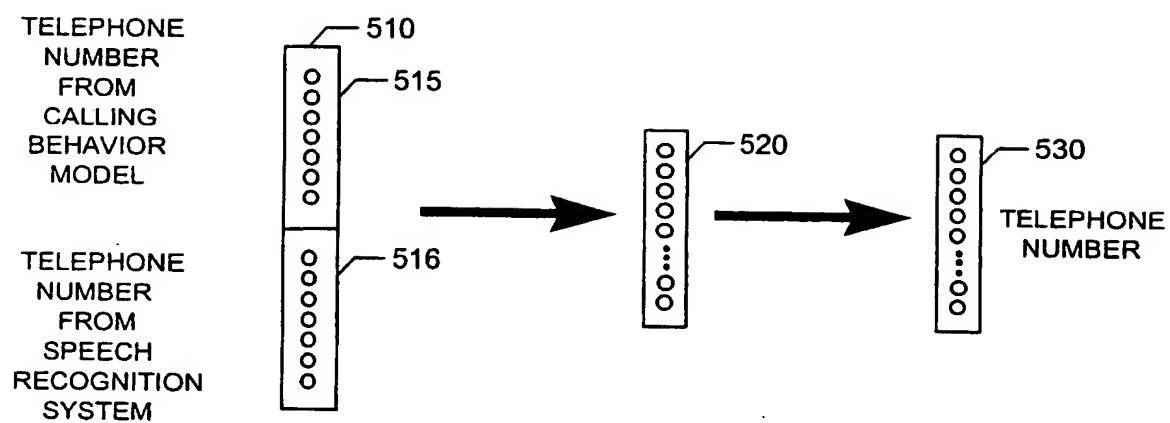
3/16

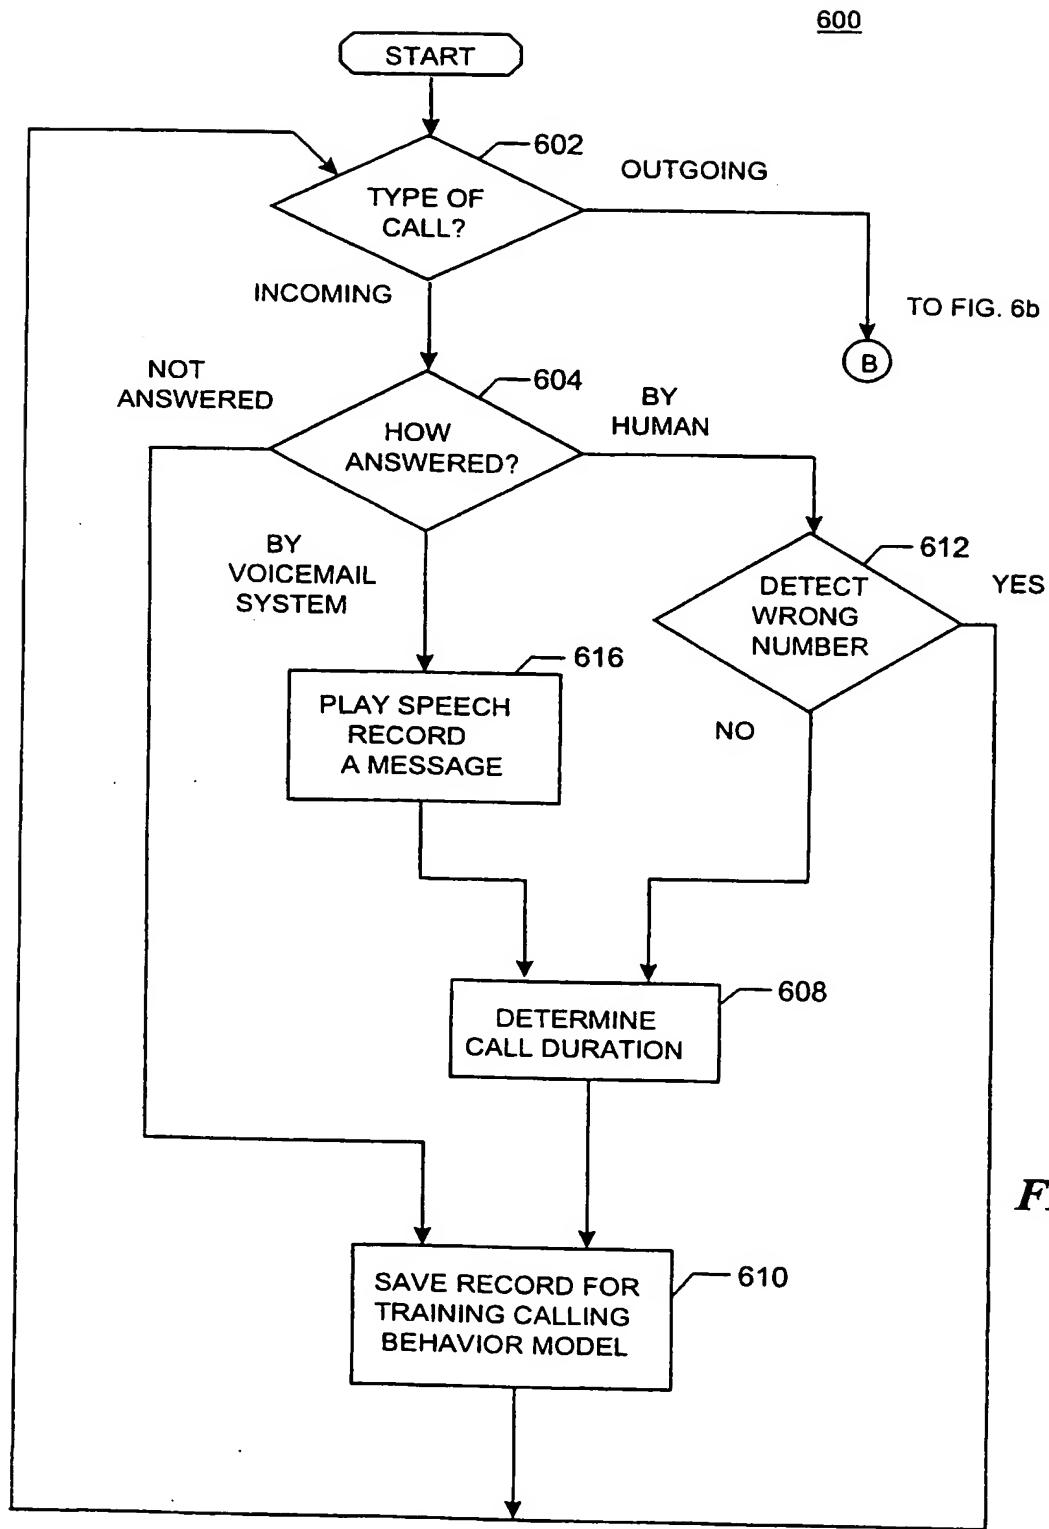
300**FIG. 3****SUBSTITUTE SHEET (RULE 26)**

4/16

400**FIG. 4**

5/16

500**FIG. 5****SUBSTITUTE SHEET (RULE 26)**



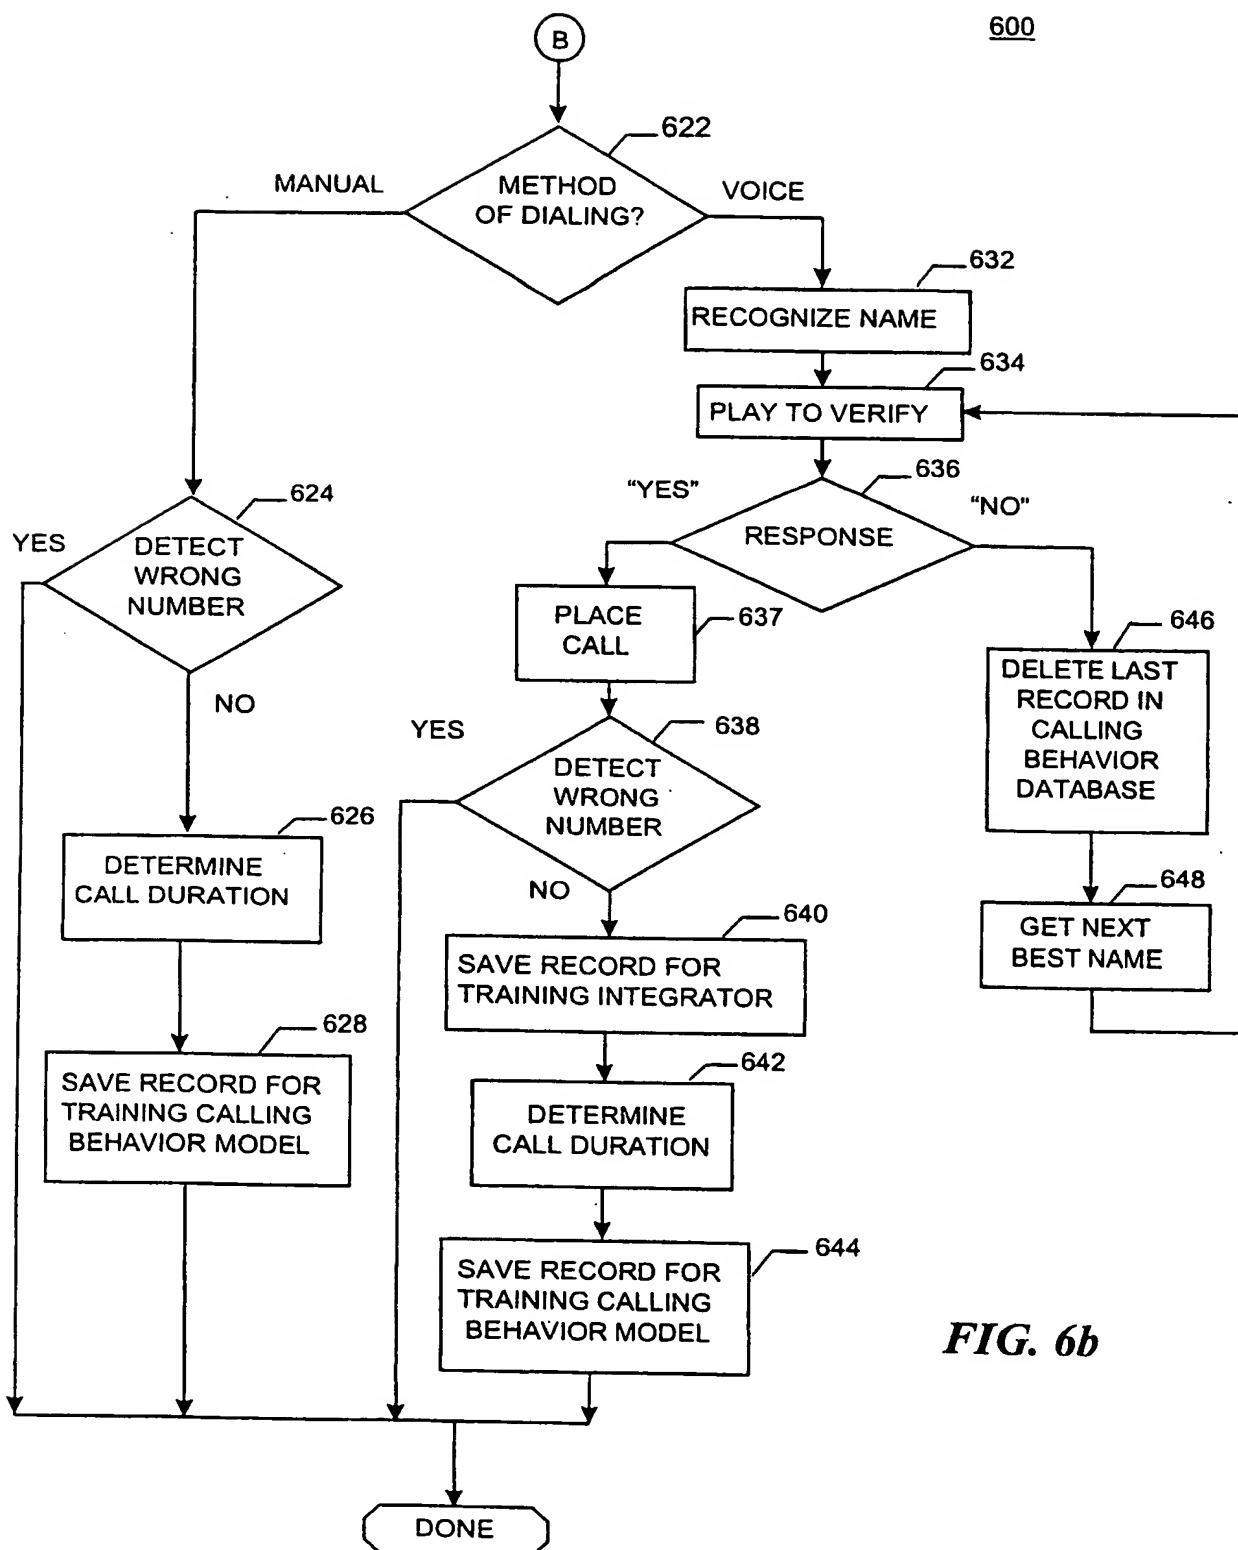


FIG. 6b

SUBSTITUTE SHEET (RULE 26)

8/16

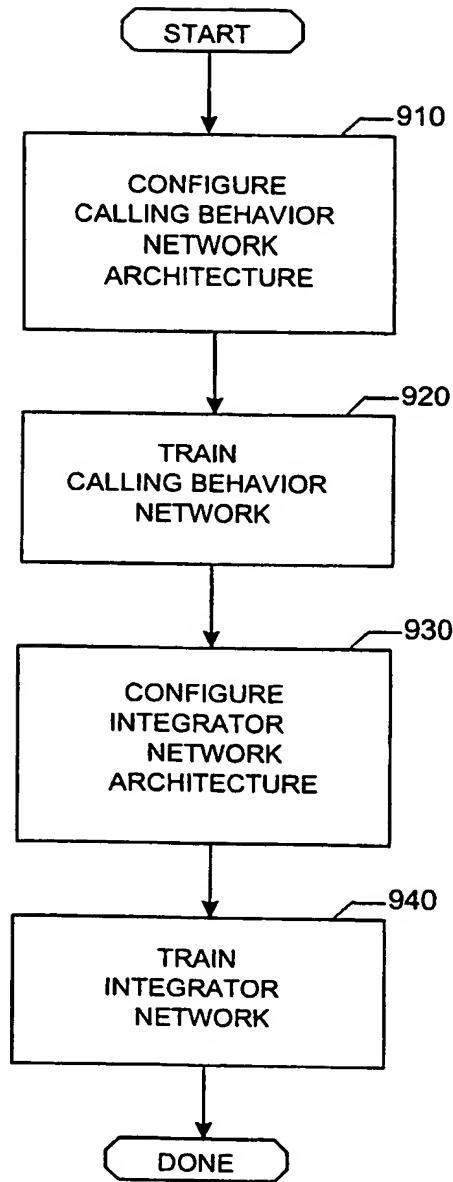
DATE OF CALL	DAY OF WEEK	TIME OF DAY	NUMBER	IN / OUT	HOW ANSWERED	DURATION OF CALL
9-22-96	2	4	408 427-8732	0	0	1

FIG. 7

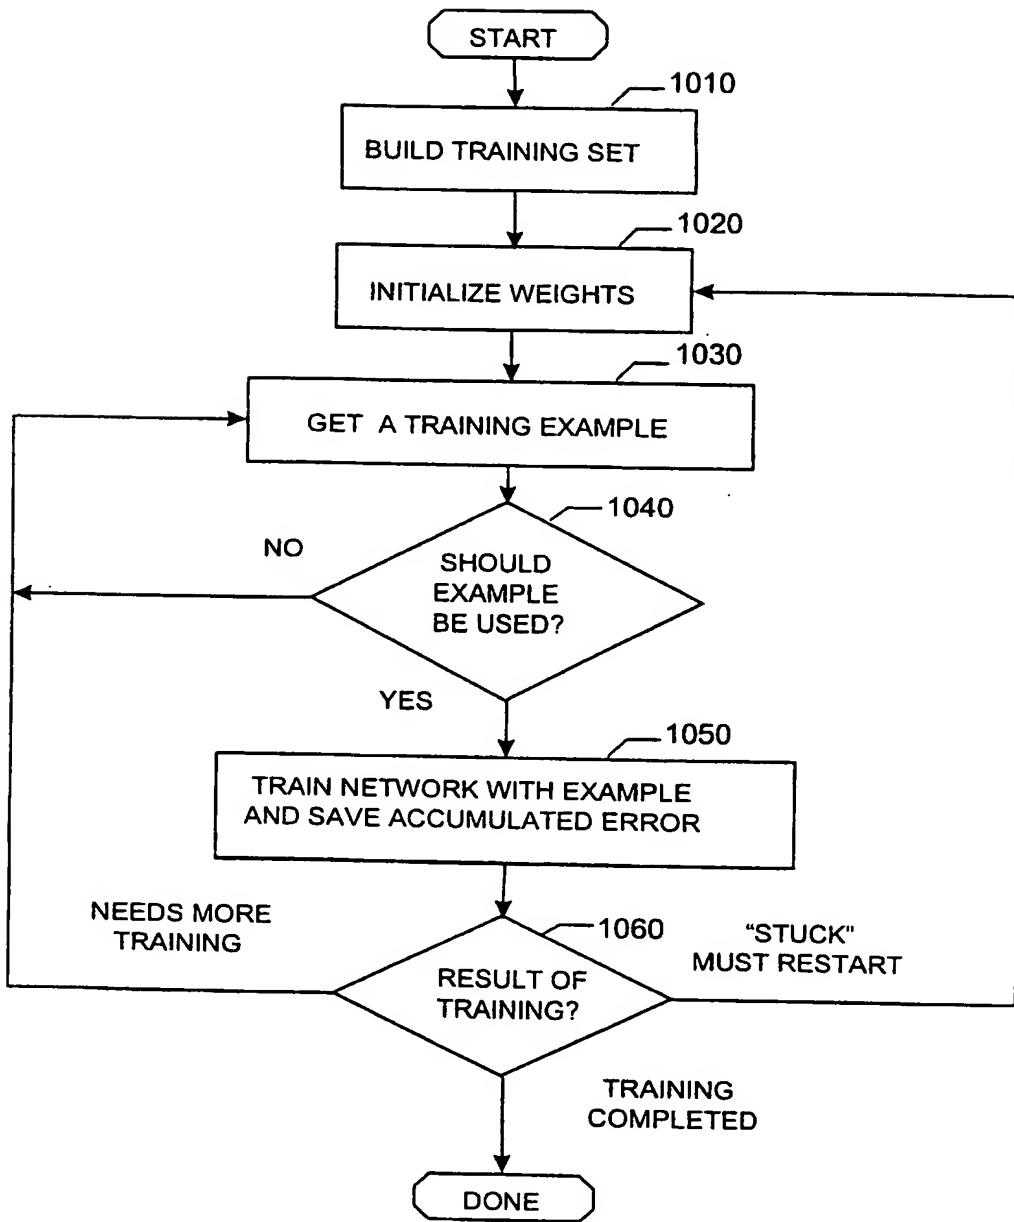
NUMBER AND MATCH QUALITY		NUMBER AND MATCH QUALITY		NUMBER AND MATCH QUALITY		CORRECT NUMBER
510 745-8224	.54	408 427-8732	.46	617 238-2189	.23	408 427-8732

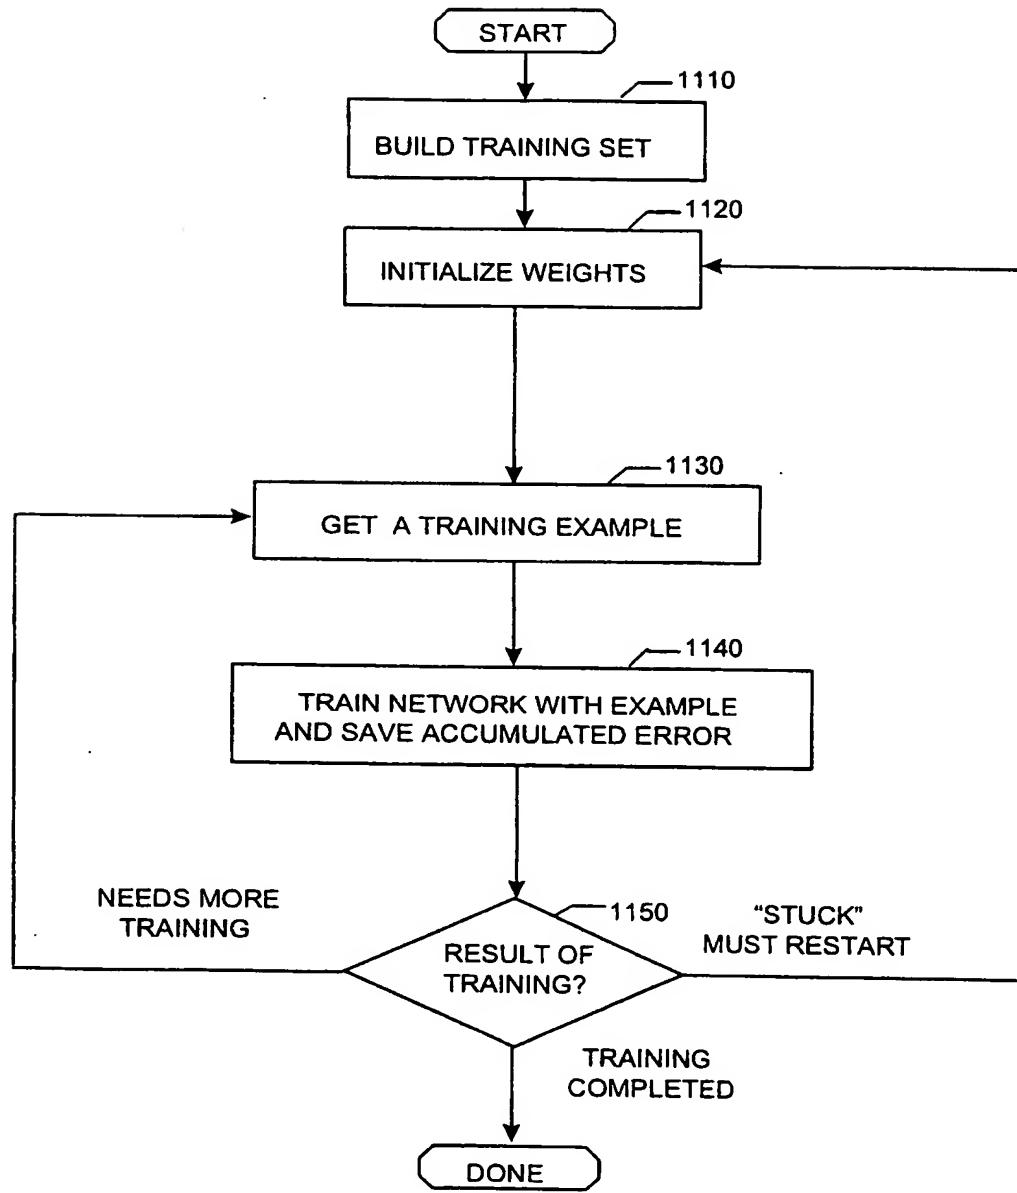
FIG. 8

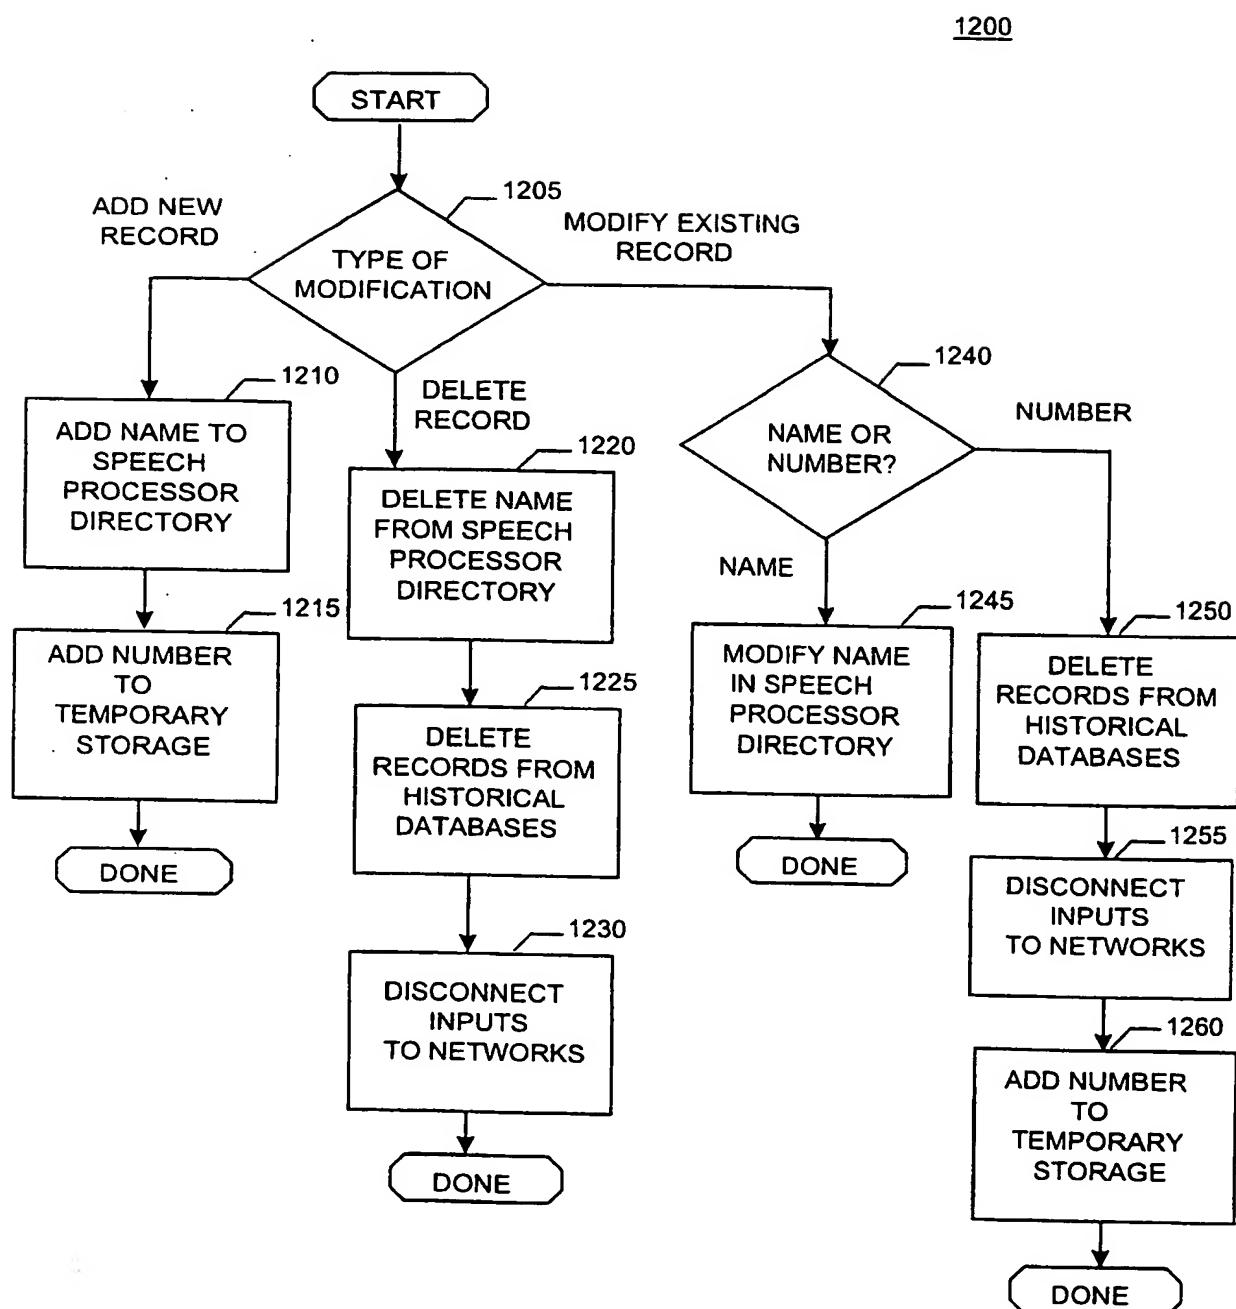
9/16

900**FIG. 9**

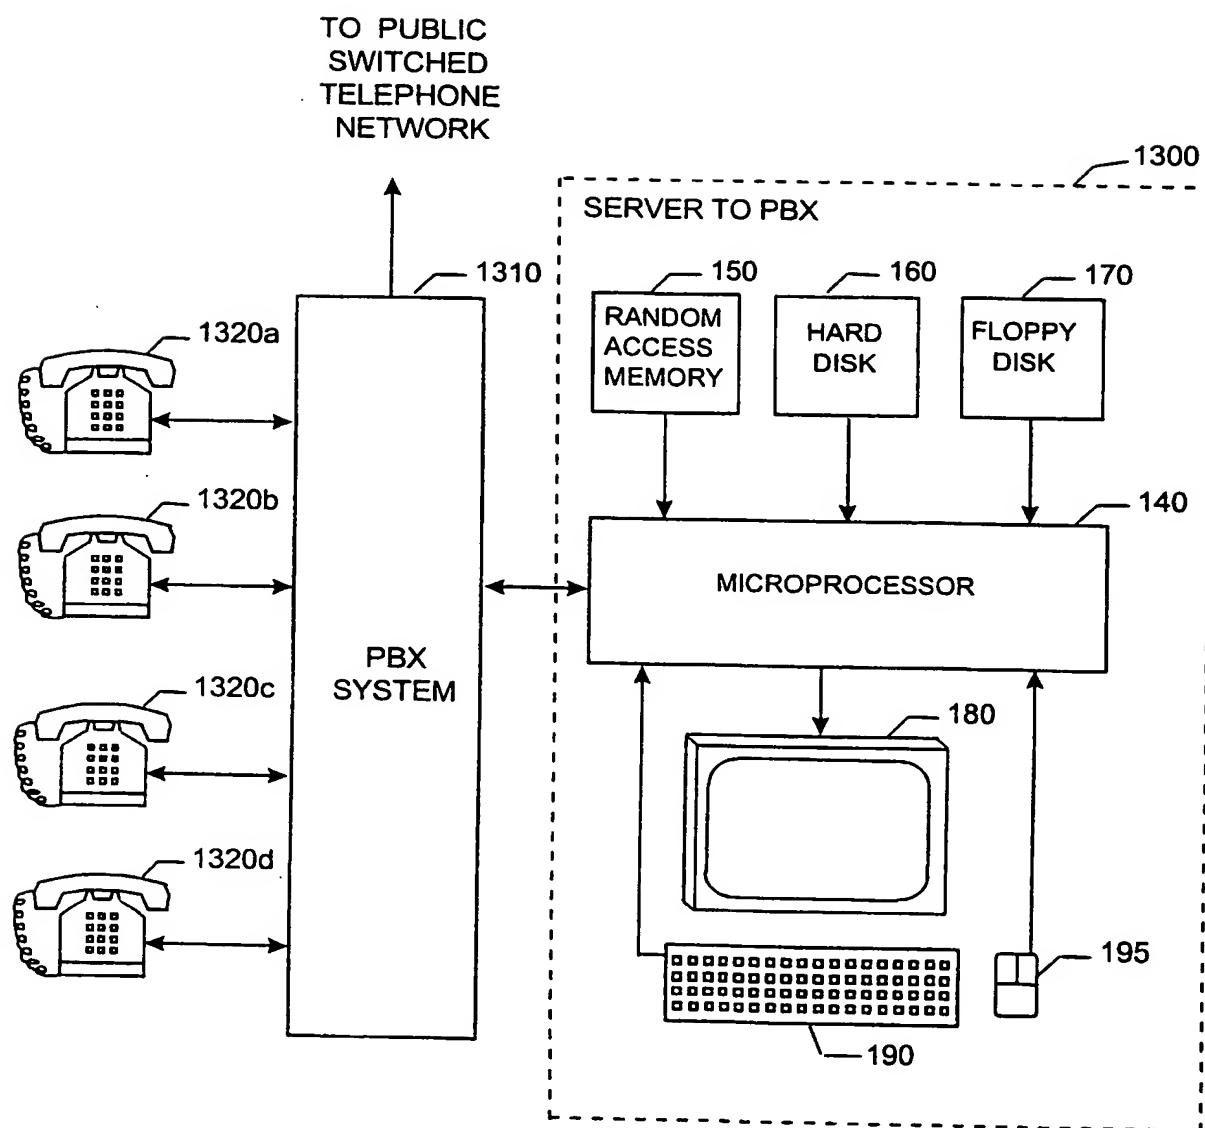
10/16

920**FIG. 10**

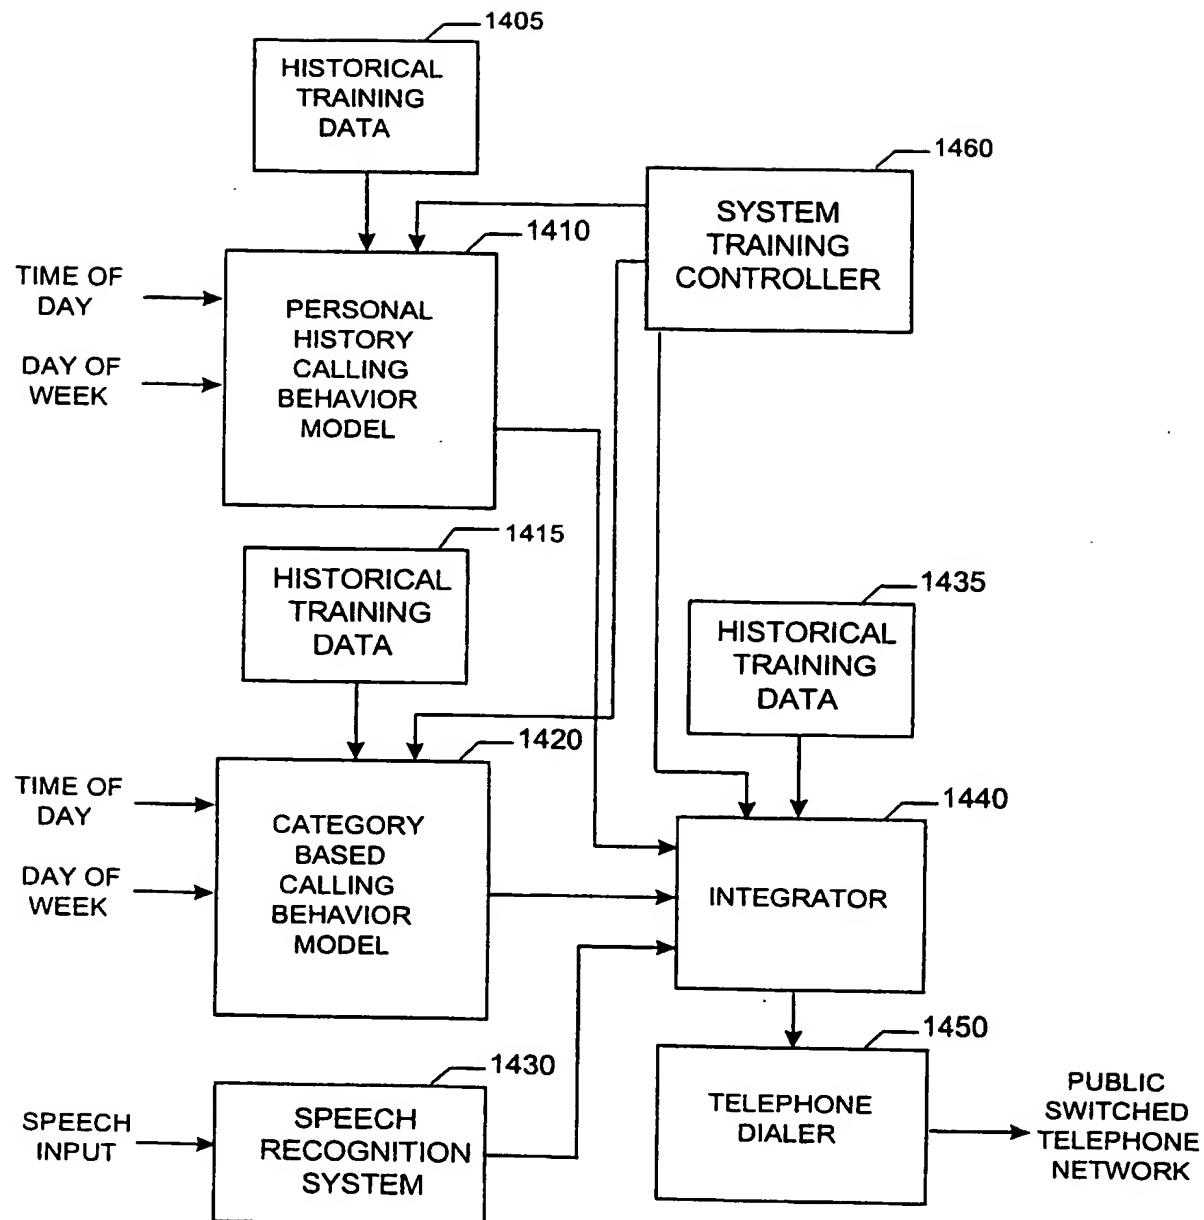
940**FIG. 11**

**FIG. 12**

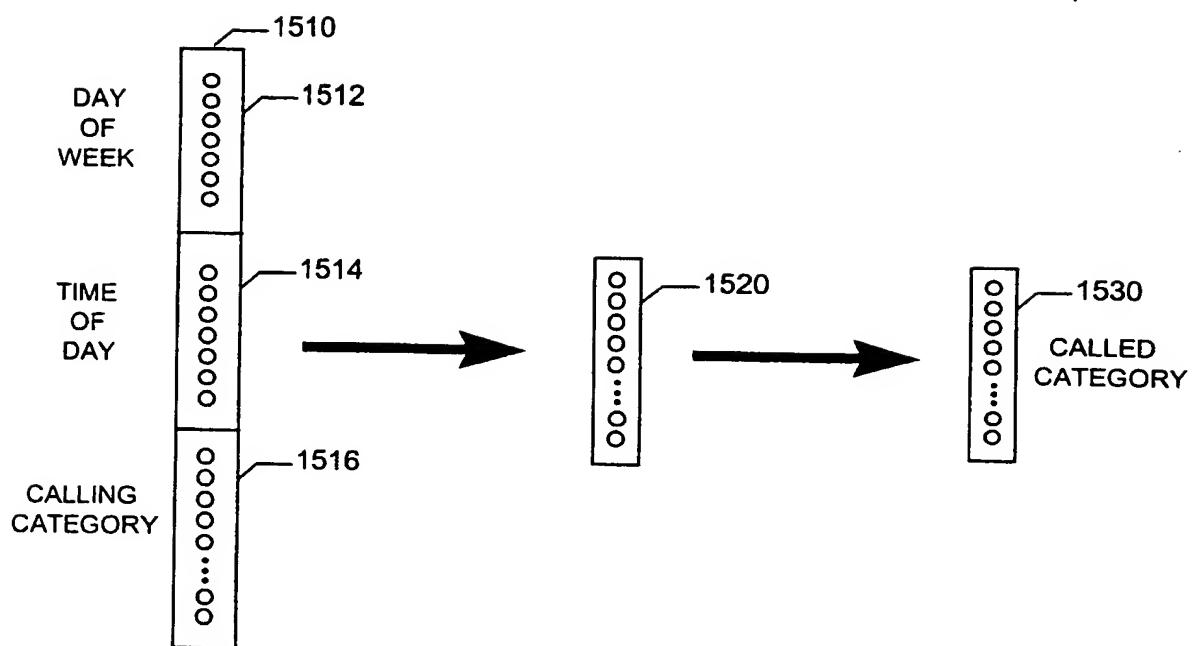
13/16

**FIG. 13****SUBSTITUTE SHEET (RULE 26)**

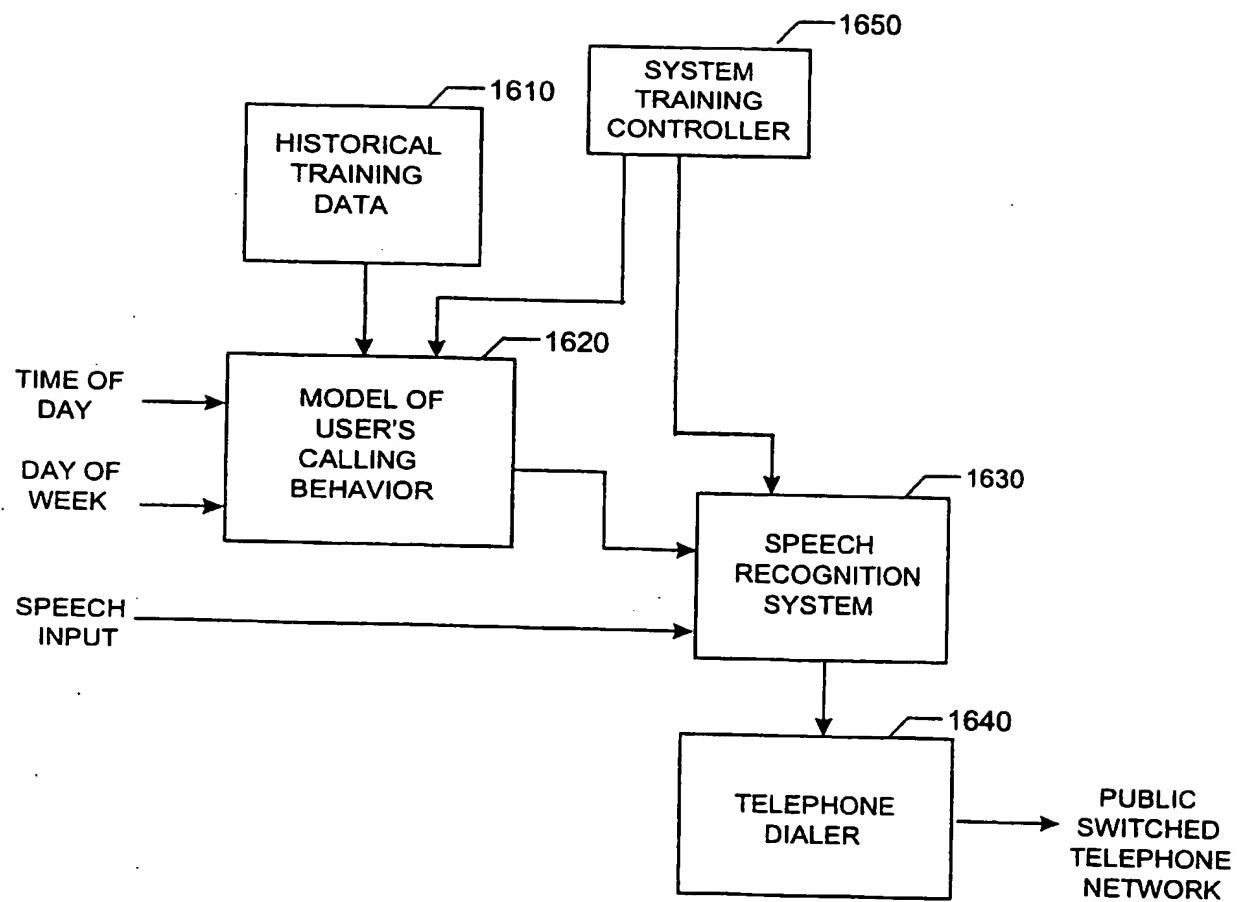
14/16

1400**FIG. 14**

15/16

1500**FIG. 15****SUBSTITUTE SHEET (RULE 26)**

16/16

1600**FIG. 16**

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 97/17623

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04M1/27

According to International Patent Classification(IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 H04M G10L H01M G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	DE 38 37 385 A (RICOH KK) 24 May 1989 see claims 1,9;10 ---	1,33,71, 72
Y	WO 87 07460 A (INNOVATIVE DEVICES) 3 December 1987 see claim 1 ---	1-9, 33-41, 71,72
Y	EP 0 269 364 A (AMERICAN TELEPHONE & TELEGRAPH) 1 June 1988 see claims 10-13 ---	1-9, 33-41, 71,72
		-/-

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

2

Date of the actual completion of the international search

Date of mailing of the international search report

27 January 1998

05/02/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Van Doremalen, J

INTERNATIONAL SEARCH REPORT

Inter nal Application No
PCT/US 97/17623

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>FANTY M ET AL: "CITY NAME RECOGNITION OVER THE TELEPHONE" PLENARY, SPECIAL, AUDIO, UNDERWATER ACOUSTICS, VLSI, NEURAL NETWORK, MINNEAPOLIS, APR. 27 - 30, 1993, vol. 1 OF 5, 27 April 1993, INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, pages 549-552, XP000399180 see abstract</p> <p>---</p>	13, 45
A	<p>PATENT ABSTRACTS OF JAPAN vol. 018, no. 218 (E-1539), 19 April 1994 & JP 06 014098 A (SHARP CORP), 21 January 1994, see abstract</p> <p>-----</p>	1-72

INTERNATIONAL SEARCH REPORT

Information on patent family members

Inter	national Application No
PCT/US 97/17623	

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
DE 3837385 A	24-05-89	JP 2023758 A		25-01-90
		JP 1295552 A		29-11-89
		JP 1295553 A		29-11-89
		JP 1251849 A		06-10-89
		JP 1123557 A		16-05-89
		GB 2212365 A, B		19-07-89
		US 4928302 A		22-05-90
WO 8707460 A	03-12-87	AU 7488287 A		22-12-87
EP 0269364 A	01-06-88	US 4862498 A		29-08-89
		AU 578861 B		03-11-88
		AU 8162987 A		02-06-88
		CA 1297201 A		10-03-92
		DE 3788786 D		24-02-94
		DE 3788786 T		05-05-94
		HK 100694 A		30-09-94
		JP 2033103 C		19-03-96
		JP 7058992 B		21-06-95
		JP 63148321 A		21-06-88
		SG 52394 A		13-01-95